

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
18 December 2003 (18.12.2003)

PCT

(10) International Publication Number
WO 03/104943 A2

(51) International Patent Classification⁷: **G06F**
(21) International Application Number: PCT/US03/18386
(22) International Filing Date: 10 June 2003 (10.06.2003)
(25) Filing Language: English
(26) Publication Language: English
(30) Priority Data:
60/388,407 11 June 2002 (11.06.2002) US
(71) Applicant and
(72) Inventor: **PANDYA, Ashish, A.** [IN/US]; 4318 Lafayette
Drive, El Dorado Hills, CA 95762 (US).

(74) Agent: **LEAL, Peter, R.**; Gray Cary Ware & Frیدنrich
LLP, Patent Department, 1755 Embarcadero Road, Palo
Alto, CA 94303 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,

CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,
MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE,
SG, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU,
ZA, ZM, ZW.

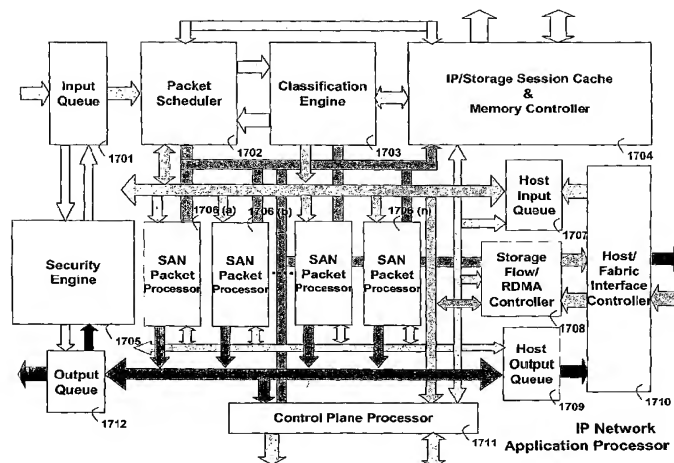
(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),
European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE,
ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO,
SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM,
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished
upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.

(54) Title: HIGH PERFORMANCE IP PROCESSOR FOR TCP/IP, RDMA AND IP STORAGE APPLICATIONS



(57) **Abstract:** An architecture provides capabilities to transport and process Internet Protocol (IP) packets from Layer (2) through full TCP/IP termination and deep packet inspection through Layer (7). A set of engines performs pass-through packet classification, policy processing and/or security processing enabling packet streaming through the architecture at nearly the full line rate. A scheduler schedules packets to packet processors for processing, based at least in part on the classification. An internal memory or local session database cache stores a TCP/IP session information database and may also store a storage information session database for a certain number of active sessions. The session information that is not in the internal memory is stored and retrieved to/from an additional memory. An application running on an initiator or target can in certain instantiations register a region of memory, which is made available to its peer(s) for access directly without substantial host intervention through RDMA data transfer.

HIGH PERFORMANCE IP PROCESSOR FOR TCP/IP, RDMA AND IP STORAGE APPLICATIONS

RELATED APPLICATIONS

Priority is claimed to Provisional Application Serial No. 60/388,407, filed on June 11, 2002.

- 5 U.S. Patent Application number not yet assigned filed on June 10, 2003 entitled High Performance IP Processor Using RDMA, U.S. Patent Application number not yet assigned filed on June 10, 2003 entitled TCP/IP Processor and Engine Using RDMA, U.S. Patent Application number not yet assigned filed on June 10, 2003 entitled IP Storage Processor and Engine Therefor Using RDMA, U.S. Patent Application number not yet assigned filed on
- 10 June 10, 2003 entitled A Memory System For A High Performance IP Processor, U.S. Patent Application number not yet assigned filed on June 10, 2003 entitled Data Processing System Using Internet Protocols and RDMA, U.S. Patent Application number not yet assigned filed on June 10, 2003 entitled High Performance IP Processor, U.S. Patent Application number not yet assigned filed on June 10, 2003 entitled Data Processing System Using Internet
- 15 Protocols, are related to the foregoing provisional application.

BACKGROUND OF THE INVENTION

This invention relates generally to storage networking semiconductors and in particular to a high performance network storage processor that is used to create Internet Protocol (IP) based storage networks.

- 20 Internet protocol (IP) is the most prevalent networking protocol deployed across various networks like local area networks (LANs), metro area networks (MANs) and wide area networks (WANs). Storage area networks (SANs) are predominantly based on Fibre Channel (FC) technology. There is a need to create IP based storage networks.

- When transporting block storage traffic on IP designed to transport data streams, the data
- 25 streams are transported using Transmission Control Protocol (TCP) that is layered to run on top of IP. TCP/IP is a reliable connection/session oriented protocol implemented in software within the operating systems. TCP/IP software stack is very slow to handle the high line rates that will be deployed in future. Currently, a 1 GHz processor based server running TCP/IP stack, with a 1Gbps network connection, would use 50-70% or more of the
- 30 processor cycles, leaving minimal cycles available for the processor to allocate to the

applications that run on the server. This overhead is not tolerable when transporting storage data over TCP/IP as well as for high performance IP networks. Hence, new hardware solutions would accelerate the TCP/IP stack to carry storage and network data traffic and be competitive to FC based solutions. In addition to the TCP protocol, other protocols such as
5 SCTP and UDP protocols can be used, as well as other protocols appropriate for transporting data streams.

SUMMARY OF THE INVENTION

I describe a high performance hardware processor that sharply reduces the TCP/IP protocol stack overhead from host processor and enables a high line rate storage and data transport
10 solution based on IP.

Traditionally, TCP/IP networking stack is implemented inside the operating system kernel as a software stack. The software TCP/IP stack implementation consumes, as mentioned above, more than 50% of the processing cycles available in a 1 GHz processor when serving a 1Gbps network. The overhead comes from various aspects of the software
15 TCP/IP stack including checksum calculation, memory buffer copy, processor interrupts on packet arrival, session establishment, session tear down and other reliable transport services. The software stack overhead becomes prohibitive at higher line rates. Similar issues occur in networks with lower line rates, like wireless networks, that use lower performance host processors. A hardware implementation can remove the overhead from
20 the host processor.

The software TCP/IP networking stack provided by the operating systems uses up a majority of the host processor cycles. TCP/IP is a reliable transport that can be run on unreliable data links. Hence, when a network packet is dropped or has errors, TCP does the retransmission of the packets. The errors in packets are detected using checksum that is
25 carried within the packet. The recipient of a TCP packet performs the checksum of the received packet and compares that to the received checksum. This is an expensive compute intensive operation performed on each packet involving each received byte in the packet. The packets between a source and destination may arrive out of order and the TCP layer performs ordering of the data stream before presenting it to the upper layers. IP
30 packets may also be fragmented based on the maximum transfer unit (MTU) of the link layer and hence the recipient is expected to de-fragment the packets. These functions result in temporarily storing the out of order packets, fragmented packets or unacknowledged packets

in memory on the network card for example. When the line rates increase to above 1Gbps, the memory size overhead and memory speed bottleneck resulting from these add significant cost to the network cards and also cause huge performance overhead. Another function that consumes a lot of processor resources is the copying of the data to/from the network card buffers, kernel buffers and the application buffers.

Microprocessors are increasingly achieving their high performance and speed using deep pipelining and super scalar architectures. Interrupting these processors on arrival of small packets will cause severe performance degradation due to context switching overhead, pipeline flushes and refilling of the pipelines. Hence interrupting the processors should be minimized to the most essential interrupts only. When the block storage traffic is transported over TCP/IP networks, these performance issues become critical, severely impacting the throughput and the latency of the storage traffic. Hence the processor intervention in the entire process of transporting storage traffic needs to be minimized for IP based storage solutions to have comparable performance and latency as other specialized network architectures like fibre channel, which are specified with a view to a hardware implementation. Emerging IP based storage standards like iSCSI, FCIP, iFCP, and others (like NFS, CIFS, DAFS, HTTP, XML, XML derivatives (such as Voice XML, EBXML, Microsoft SOAP and others), SGML, and HTML formats) encapsulate the storage and data traffic in TCP/IP segments. However, there usually isn't alignment relationship between the TCP segments and the protocol data units that are encapsulated by TCP packets. This becomes an issue when the packets arrive out of order, which is a very frequent event in today's networks. The storage and data blocks cannot be extracted from the out of order packets for use until the intermediate packets in the stream arrive which will cause the network adapters to store these packets in the memory, retrieve them and order them when the intermediate packets arrive. This can be expensive from the size of the memory storage required and also the performance that the memory subsystem is expected to support, particularly at line rates above 1Gbps. This overhead can be removed if each TCP segment can uniquely identify the protocol data unit and its sequence. This can allow the packets to be directly transferred to their end memory location in the host system. Host processor intervention should also be minimized in the transfer of large blocks of data that may be transferred to the storage subsystems or being shared with other processors in a clustering environment or other client server environment. The processor should be interrupted only on storage command boundaries to minimize the impact.

The IP processor set forth herein eliminates or sharply reduces the effect of various issues outlined above through innovative architectural features and the design. The described processor architecture provides features to terminate the TCP traffic carrying the storage and data payload thereby eliminating or sharply reducing the TCP/IP networking stack overhead on the host processor, resulting in packet streaming architecture that allows packets to pass through from input to output with minimal latency. To enable high line rate storage or data traffic being carried over IP requires maintaining the transmission control block information for various connections (sessions) that are traditionally maintained by host kernel or driver software. As used in this patent, the term "IP session" means a session for a session oriented protocol that runs on IP. Examples are TCP/IP, SCTP/IP, and the like. Accessing session information for each packet adds significant processing overhead. The described architecture creates a high performance memory subsystem that significantly reduces this overhead. The architecture of the processor provides capabilities for intelligent flow control that minimizes interrupts to the host processor primarily at the command or data transfer completion boundary.

Today, no TCP/IP processor is offered with security.

The described processor architecture also provides integrated security features. When the storage traffic is carried on a network from the server to the storage arrays in a SAN or other storage system, it is exposed to various security vulnerabilities that a direct attached storage system does not have to deal with. This processor allows for in stream encryption and decryption of the storage traffic thereby allowing high line rates and at the same time offering confidentiality of the storage data traffic.

Classification of network traffic is another task that consumes up to half of the processing cycles available on packet processors leaving few cycles for deep packet inspection and processing. IP based storage traffic by the nature of the protocol requires high speed low latency deep packet processing. The described IP processor significantly reduces the classification overhead by providing a programmable classification engine.

Tremendous growth in the storage capacity and storage networks have created storage area management as a major cost item for IT departments. Policy based storage management is required to contain management costs. The described programmable classification engine allows deployment of storage policies that can be enforced on packet, transaction, flow and

command boundaries. This will have significant improvement in storage area management costs.

The programmable IP processor architecture also offers enough headroom to allow customer specific applications to be deployed. These applications may belong to multiple categories e.g. network management, storage firewall or other security capabilities, bandwidth management, quality of service, virtualization, performance monitoring, zoning, LUN masking and the like.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates a layered SCSI architecture and interaction between respective layers located between initiator and target systems.

Fig. 2 illustrates the layered SCSI architecture with iSCSI and TCP/IP based transport between initiator and target systems.

Fig. 3 illustrates an OSI stack comparison of software based TCP/IP stack with hardware - oriented protocols like Fibre channel.

Fig. 4 illustrates an OSI stack with a hardware based TCP/IP implementation for providing performance parity with the other non-IP hardware oriented protocols.

Fig. 5 illustrates a host software stack illustrating operating system layers implementing networking and storage stacks.

Fig. 6 illustrates software TCP stack data transfers.

Fig. 7 illustrates remote direct memory access data transfers using TCP/IP offload from the host processor as described in this patent.

Fig. 8 illustrates host software SCSI storage stack layers for transporting block storage data over IP networks.

Fig. 9 illustrates certain iSCSI storage network layer stack details of an embodiment of the invention.

Fig. 10 illustrates TCP/IP network stack functional details of an embodiment of the invention.

Fig. 11 illustrates an iSCSI storage data flow through various elements of an embodiment of the invention.

Fig. 12 illustrates iSCSI storage data structures useful in the invention.

- 5 Fig. 13 illustrates a TCP/IP Transmission Control Block data structure for a session database entry useful in an embodiment of the invention.

Fig. 14 illustrates an iSCSI session database structure useful in an embodiment of the invention.

Fig. 15 illustrates iSCSI session memory structure useful in an embodiment of the invention.

- 10 Fig. 16 illustrates a high-level architectural block diagram of an IP network application processor useful in an embodiment of the invention.

Fig. 17 illustrates a detailed view of the architectural block diagram of the IP network application processor of Fig. 16.

Fig. 18 illustrates an input queue and controller for one embodiment of the IP processor.

- 15 Fig. 19 illustrates a packet scheduler, sequencer and load balancer useful in one embodiment of the IP processor.

Fig. 20 illustrates a packet classification engine, including a policy engine block of one embodiment of the IP storage processor.

- 20 Fig. 21 broadly illustrates an embodiment of the SAN packet processor block of one embodiment of an IP processor at a high-level.

Fig. 22 illustrates an embodiment of the SAN packet processor block of the described IP processor in further detail.

Fig. 23 illustrates an embodiment of the programmable TCP/IP processor engine which can be used as part of the described SAN packet processor.

Fig. 24 illustrates an embodiment of the programmable IP Storage processor engine which can be used as part of the described SAN packet processor.

Fig. 25 illustrates an embodiment of an output queue block of the programmable IP processor of Fig. 17.

5 Fig. 26 illustrates an embodiment of the storage flow controller and RDMA controller.

Fig. 27 illustrates an embodiment of the host interface controller block of the IP processor useful in an embodiment of the invention.

Fig. 28 illustrates an embodiment of the security engine.

10 Fig. 29 illustrates an embodiment of a memory and controller useful in the described processor.

Fig. 30 illustrates a data structure useable in an embodiment of the described classification engine.

Fig. 31 illustrates a storage read flow between initiator and target.

15 Fig. 32 illustrates a read data packet flow through pipeline stages of the described processor.

Fig. 33 illustrates a storage write operation flow between initiator and target.

Fig. 34 illustrates a write data packet flow through pipeline stages of the described processor.

20 Fig. 35 illustrates a storage read flow between initiator and target using the remote DMA (RDMA) capability between initiator and target.

Fig. 36 illustrates a read data packet flow between initiator and target using RDMA through pipeline stages of the described processor.

Fig. 37 illustrates a storage write flow between initiator and target using RDMA capability.

Fig. 38 illustrates a write data packet flow using RDMA through pipeline stages of the described processor.

Fig. 39 illustrates an initiator command flow in more detail through pipeline stages of the described processor.

- 5 Fig. 40 illustrates a read packet data flow through pipeline stages of the described processor in more detail.

Fig. 41 illustrates a write data flow through pipeline stages of the described processor in more detail.

- 10 Fig. 42 illustrates a read data packet flow when the packet is in cipher text or is otherwise a secure packet through pipeline stages of the described processor.

Fig. 43 illustrates a write data packet flow when the packet is in cipher text or is otherwise a secure packet through pipeline stages of the described processor of one embodiment of the invention.

- 15 Fig. 44 illustrates a RDMA buffer advertisement flow through pipeline stages of the described processor.

Fig. 45 illustrates a RDMA write flow through pipeline stages of the described processor in more detail.

Fig. 46 illustrates a RDMA Read data flow through pipeline stages of the described processor in more detail.

- 20 Fig. 47 illustrates steps of a session creation flow through pipeline stages of the described processor.

Fig. 48 illustrates steps of a session tear down flow through pipeline stages of the described processor.

- 25 Fig. 49 illustrates a session creation and session teardown steps from a target perspective through pipeline stages of the described processor.

Fig. 50 illustrates an R2T command flow in a target subsystem through pipeline stages of the described processor.

Fig. 51 illustrates a write data flow in a target subsystem through pipeline stages of the described processor.

- 5 Fig. 52 illustrates a target read data flow through the pipeline stages of the described processor.

DESCRIPTION

- I provide a new high performance and low latency way of implementing a TCP/IP stack in hardware to relieve the host processor of the severe performance impact of a software
10 TCP/IP stack. This hardware TCP/IP stack is then interfaced with additional processing elements to enable high performance and low latency IP based storage applications.

- This can be implemented in a variety of forms to provide benefits of TCP/IP termination, high performance and low latency IP storage capabilities, remote DMA (RDMA) capabilities, security capabilities, programmable classification and policy processing features and the
15 like. Following are some of the embodiments that can implement this:

Server

- The described architecture may be embodied in a high performance server environment providing hardware based TCP/IP functions that relieve the host server processor or processors of TCP/IP software and performance overhead. The IP processor may be a
20 companion processor to a server chipset, providing the high performance networking interface with hardware TCP/IP. Servers can be in various form factors like blade servers, appliance servers, file servers, thin servers, clustered servers, database server, game server, grid computing server, VOIP server, wireless gateway server, security server, network attached storage server or traditional servers. The current embodiment would allow
25 creation of a high performance network interface on the server motherboard.

Companion Processor to a server Chipset

The server environment may also leverage the high performance IP storage processing capability of the described processor, besides high performance TCP/IP and/or RDMA capabilities. In such an embodiment the processor may be a companion processor to a

server chipset providing high performance network storage I/O capability besides the TCP/IP offloading from the server processor. This embodiment would allow creation of high performance IP based network storage I/O on the motherboard. In other words it would enable IP SAN on the motherboard.

5 **Storage System Chipsets**

The processor may also be used as a companion of a chipset in a storage system, which may be a storage array (or some other appropriate storage system or subsystem) controller, which performs the storage data server functionality in a storage networking environment. The processor would provide IP network storage capability to the storage array controller to
10 network in an IP based SAN. The configuration may be similar to that in a server environment, with additional capabilities in the system to access the storage arrays and provide other storage-centric functionality.

Server/Storage Host Adapter Card

The IP processor may also be embedded in a server host adapter card providing high speed
15 TCP/IP networking. The same adapter card may also be able to offer high speed network storage capability for IP based storage networks. The adapter card may be used in traditional servers and may also be used as blades in a blade server configuration. The processor may also be used in adapters in a storage array (or other storage system or subsystem) front end providing IP based storage networking capabilities.

20 **Processor Chipset Component**

The TCP/IP processor may be embodied inside a processor chipset, providing the TCP/IP offloading capability. Such a configuration may be used in the high end servers, workstations or high performance personal computers that interface with high speed networks. Such an embodiment could also include IP storage or RDMA capabilities or
25 combination of this invention to provide IP based storage networking and/or TCP/IP with RDMA capability embedded in the chipset. The usage of multiple capabilities of the described architecture can be made independent of using other capabilities in this or other embodiments, as a trade-off of feature requirements, development timeline and cost, silicon die cost and the like.

Storage or SAN System or Subsystem Switching Line Cards

The IP processor may also be used to create high performance, low latency IP SAN switching system (or other storage system or subsystem) line cards. The processor may be used as the main processor terminating and originating IP-based storage traffic to/from the line card. This processor would work with the switching system fabric controller, which may act like a host, to transport the terminated storage traffic, based on their IP destination, to the appropriate switch line card as determined by the forwarding information base present in the switch system. Such a switching system may support purely IP based networking or may support multi-protocol support, allow interfacing with IP based SAN along with other data center SAN fabrics like Fibre channel. A very similar configuration could exist inside a gateway controller system, that terminates IP storage traffic from LAN or WAN and originates new sessions to carry the storage traffic into a SAN, which may be IP based SAN or more likely a SAN built from other fabrics inside a data center like Fibre channel. The processor could also be embodied in a SAN gateway controller.

15 Storage Appliance

Storage networks management costs are increasing rapidly. The ability to manage the significant growth in the networks and the storage capacity would require creating special appliances which would be providing the storage area management functionality. The described management appliances for high performance IP based SAN, would implement my high performance IP processor, to be able to perform its functions on the storage traffic transported inside TCP/IP packets. These systems would require a high performance processor to do deep packet inspection and extract the storage payload in the IP traffic to provide policy based management and enforcement functions. The security, programmable classification and policy engines along with the high speed TCP/IP and IP storage engines described would enable these appliances and other embodiments described in this patent to perform deep packet inspection and classification and apply the policies that are necessary on a packet by packet basis at high line rates at low latency. Further these capabilities can enable creating storage management appliances that can perform their functions like virtualization, policy based management, security enforcement, access control, intrusion detection, bandwidth management, traffic shaping, quality of service, anti-spam, virus detection, encryption, decryption, LUN masking, zoning, link aggregation and the like in-band to the storage area network traffic. Similar policy based management, and security

operations or functionality may also be supported inside the other embodiments described in this patent.

Clustered Environments

Server systems are used in a clustered environment to increase the system performance and scalability for applications like clustered data bases and the like. The applications
5 running on high performance cluster servers require ability to share data at high speeds for inter-process communication. Transporting this inter-process communication traffic on a traditional software TCP/IP network between cluster processors suffers from severe performance overhead. Hence, specialized fabrics like Fibre channel have been used in
10 such configurations. However, a TCP/IP based fabric which can allow direct memory access between the communicating processes' memory, can be used by applications that operate on any TCP/IP network without being changed to specialized fabrics like fibre channel. The described IP processor with its high performance TCP/IP processing capability and the RDMA features, can be embodied in a cluster server environment to provide the benefits of
15 high performance and low latency direct memory to memory data transfers. This embodiment may also be used to create global clustering and can also be used to enable data transfers in grid computers and grid networks.

Additional Embodiments

The processor architecture can be partially implemented in software and partially in
20 hardware. The performance needs and cost implications can drive trade-offs for hardware and software partitioning of the overall system architecture of this invention. It is also possible to implement this architecture as a combination of chip sets along with the hardware and software partitioning or independent of the partitioning. For example the security processor and the classification engines could be on separate chips and provide
25 similar functions. This can result in lower silicon cost of the IP processor including the development and manufacturing cost, but it may in some instances increase the part count in the system and may increase the footprint and the total solution cost. Security and classification engines could be separate chips as well. As used herein, a chip set may mean a multiple-chip chip set, or a chip set that includes only a single chip, depending on the
30 application.

The storage flow controller and the queues could be maintained in software on the host or may become part of another chip in the chipset. Hence, multiple ways of partitioning this

architecture are feasible to accomplish the high performance IP based storage and TCP/IP offload applications that will be required with the coming high performance processors in the future. The storage engine description has been given with respect to iSCSI, however, with TCP/IP and storage engine programmability, classifier programmability and the storage flow controller along with the control processor, other IP storage protocols like iFCP, FCIP and others can be implemented with the appropriate firmware. iSCSI operations may also be IP Storage operations. The high performance IP processor core may be coupled with multiple input output ports of lower line rates, matching the total throughput to create multi-port IP processor embodiment as well.

- 10 It is feasible to use this architecture for high performance TCP/IP offloading from the main processor without using the storage engines. This can result in a silicon and system solution for next generation high performance networks for the data and telecom applications. The TCP/IP engine can be augmented with application specific packet accelerators and leverage the core architecture to derive new flavors of this processor. It is possible to change the storage engine with another application specific accelerator like a firewall engine or a route look-up engine or a telecom/network acceleration engine, along with the other capabilities of this invention and target this processor architecture for telecom/networking and other applications.

Detailed Description

- 20 Storage costs and demand have been increasing at a rapid pace over the last several years. This is expected to grow at the same rate in the foreseeable future. With the advent of e-business, availability of the data at any time and anywhere irrespective of the server or system downtime is critical. This is driving a strong need to move the server attached storage onto a network to provide storage consolidation, availability of data and ease of management of the data. The storage area networks (SANs) are today predominantly based on Fibre Channel technology, that provide various benefits like low latency and high performance with its hardware oriented stacks compared to TCP/IP technology.

- 30 Some system transport block storage traffic on IP designed to transport data streams. The data streams are transported using Transmission Control Protocol (TCP) that is layered to run on top of IP. TCP/IP is a reliable connection oriented protocol implemented in software within the operating systems. A TCP/IP software stack is slow to handle the high line rates

that will be deployed in the future. New hardware solutions will accelerate the TCP/IP stack to carry storage and network traffic and be competitive to FC based solutions.

The prevalent storage protocol in high performance servers, workstations and storage controllers and arrays is SCSI protocol which has been around for 20 years. SCSI
5 architecture is built as layered protocol architecture. Fig. 1 illustrates the various SCSI architecture layers within an initiator, block 101, and target subsystems, block 102. As used in patent, the terms "initiator" and "target" mean a data processing apparatus, or a subsystem or system including them. The terms "initiator" and "target" can also mean a client or a server or a peer. Likewise, the term "peer" can mean a peer data processing
10 apparatus, or a subsystem or system thereof. A "remote peer" can be a peer located across the world or across the room.

The initiator and target subsystems in Fig. 1 interact with each other using the SCSI application protocol layer, block 103, which is used to provide a client-server request and response transactions. It also provides device service request and response between the
15 initiator and the target mass storage device which may take many forms like a disk arrays, tape drives, and the like. Traditionally, the target and initiator are interconnected using the SCSI bus architecture carrying the SCSI protocol, block 104. The SCSI protocol layer is the transport layer that allows the client and the server to interact with each other using the SCSI application protocol. The transport layer must present the same semantics to the upper
20 layer so that the upper layer protocols and application can stay transport protocol independent.

Fig. 2 illustrates the SCSI application layer on top of IP based transport layers. An IETF standards track protocol, iSCSI (SCSI over IP) is an attempt to provide IP based storage transport protocol. There are other similar attempts including FCIP (FC encapsulated in IP),
25 iFCP(FC over IP) and others. Many of these protocols layer on top of TCP/IP as the transport mechanism, in a manner similar to that illustrated in Fig. 2. As illustrated in Fig. 2, the iSCSI protocol services layer, block 204, provides the layered interface to the SCSI application layer, block 203. iSCSI carries SCSI commands and data as iSCSI protocol data units (PDUs) as defined by the standard. These protocol data units then can be transported
30 over the network using TCP/IP, block 205, or the like. The standard does not specify the means of implementing the underlying transport that carries iSCSI PDUs. Fig. 2 illustrates iSCSI layered on TCP/IP which provides the transport for the iSCSI PDUs.

The IP based storage protocol like iSCSI can be layered in software on top of a software based TCP/IP stack. However, such an implementation would suffer serious performance penalties arising from software TCP/IP and the storage protocol layered on top of that. Such an implementation would severely impact the performance of the host processor and may
5 make the processor unusable for any other tasks at line rates above 1Gbps. Hence, we would implement the TCP/IP stack in hardware, relieving the host processor, on which the storage protocol can be built. The storage protocol, like iSCSI, can be built in software running on the host processor or may, as described in this patent, be accelerated using hardware implementation. A software iSCSI stack will present many interrupts to the host
10 processor to extract PDUs from received TCP segments to be able to act on them. Such an implementation will suffer severe performance penalties for reasons similar to those for which a software based TCP stack would. The described processor provides a high performance and low latency architecture to transport Storage protocol on a TCP/IP based network that eliminates or greatly reduces the performance penalty on the host processor,
15 and the resulting latency impact.

Fig. 3 illustrates a comparison of the TCP/IP stack to Fibre channel as referenced to the OSI networking stack. The TCP/IP stack, block 303, as discussed earlier in the Summary of the Invention section of this patent, has performance problems resulting from the software implementation on the hosts. Compared to that, specialized networking protocols like Fibre
20 channel, block 304, and others are designed to be implemented in hardware. The hardware implementation allows the networking solutions to be higher performance than the IP based solution. However, the ubiquitous nature of IP and the familiarity of IP from the IT users' and developers' perspective makes IP more suitable for wide spread deployment. This can be accomplished if the performance penalties resulting from TCP/IP are reduced to be
25 equivalent to those of the other competing specialized protocols. Fig. 4 illustrates a protocol level layering in hardware and software that is used for TCP/IP, block 403, to become competitive to the other illustrated specialized protocols.

Fig. 5 illustrates a host operating system stack using a hardware based TCP/IP and storage protocol implementation of this patent. The protocol is implemented such that it can be
30 introduced into the host operating system stack, block 513, such that the operating system layers above it are unchanged. This allows the SCSI application protocols to operate without any change. The driver layer, block 515, and the stack underneath for IP based

storage interface, block 501, will represent a similar interface as a non-networked SCSI interface, blocks 506 and 503 or Fibre Channel interface, block 502.

Fig. 6 illustrates the data transfers involved in a software TCP/IP stack. Such an implementation of the TCP/IP stack carries huge performance penalties from memory copy of the data transfers. The figure illustrates data transfer between client and server networking stacks. User level application buffers, block 601, that need to be transported from the client to the server or vice versa, go through the various levels of data transfers shown. The user application buffers on the source get copied into the OS kernel space buffers, block 602. This data then gets copied to the network driver buffers, block 603, from where it gets DMA-transferred to the network interface card (NIC) or the host bus adapter (HBA) buffers, block 604. The buffer copy operations involve the host processor and use up valuable processor cycles. Further, the data being transferred goes through checksum calculations on the host using up additional computing cycles from the host. The data movement into and out of the system memory on the host multiple times creates a memory bandwidth bottleneck as well. The data transferred to the NIC/HBA is then sent on to the network, block 609, and reaches the destination system. At the destination system the data packet traverses through the software networking stack in the opposite direction as the host though following similar buffer copies and checksum operations. Such implementation of TCP/IP stack is very inefficient for block storage data transfers and for clustering applications where a large amount of data may be transferred between the source and the destination.

Fig. 7 illustrates the networking stack in an initiator and in a target with features that allow remote direct memory access (RDMA) features of the architecture described in this patent. The following can be called an RDMA capability or an RDMA mechanism or an RDMA function. In such a system the application running on the initiator or target registers a region of memory, block 702, which is made available to its peer(s) for access directly from the NIC/HBA without substantial host intervention. These applications would also let their peer(s) know about the memory regions being available for RDMA, block 708. Once both peers of the communication are ready to use the RDMA mechanism, the data transfer from RDMA regions can happen with essentially zero copy overhead from the source to the destination without substantial host intervention if NIC/HBA hardware in the peers implement RDMA capability. The source, or initiator, would inform its peer of its desire to read or write

specific RDMA enabled buffers and then let the destination or target, push or pull the data to/from its RDMA buffers. The initiator and the target NIC/HBA would then transport the data using the TCP/IP hardware implementation described in this patent, RMDA 703, TCP/IP offload 704, RMDA 708 and TCP/IP offload 709, between each other without substantial intervention of the host processors, thereby significantly reducing the processor overhead. This mechanism would significantly reduce the TCP/IP processing overhead on the host processor and eliminate the need for multiple buffer copies for the data transfer illustrated in Fig. 6. RDMA enabled systems would thus allow the system, whether fast or slow, to perform the data transfer without creating a performance bottleneck for its peer. RDMA capability implemented in this processor in storage over IP solution eliminates host intervention except usually at the data transfer start and termination. This relieves the host processors in both target and initiator systems to perform useful tasks without being interrupted at each packet arrival or transfer. RDMA implementation also allows the system to be secure and prevent unauthorized access. This is accomplished by registering the exported memory regions with the HBA/NIC with their access control keys along with the region IDs. The HBA/NIC performs the address translation of the memory region request from the remote host to the RDMA buffer, performs security operations such as security key verification and then allows the data transfer. This processing is performed off the host processor in the processor of this invention residing on the HBA/NIC or as a companion processor to the host processor on the motherboard, for example. This capability can also be used for large data transfers for server clustering applications as well as client server applications. Real time media applications transferring large amounts of data between a source or initiator and a destination or target can benefit from this.

Fig. 8 illustrates the host file system and SCSI stack implemented in software. As indicated earlier the IP based storage stack, blocks 805, 806, 807, 808 and 809, should represent a consistent interface to the SCSI layers, blocks 803 and 804, as that provided by SCSI transport layer, block 811, or Fibre channel transport, block 810. This figure illustrates high level requirements that are imposed on the IP based storage implementation from a system level, besides those imposed by various issues of IP which is not designed to transport performance sensitive block data.

Fig. 9 illustrates the iSCSI stack in more detail from that illustrated in Fig. 8. The iSCSI stack blocks 805 through 809, should provide an OS defined driver interface level functionality to

the SCSI command consolidation layer blocks 803 & 804, such that the behavior of this layer and other layers on top of it are unchanged. Fig. 9 illustrates a set of functions that would be implemented to provide IP storage capabilities. The functions that provide the iSCSI functionality are grouped into related sets of functions, although there can be many variations of these as any person skilled in this area would appreciate. There are a set of functions that are required to meet the standard (e.g. target and initiator login and logout) functions, block 916, connection establishment and teardown functions, block 905. The figure illustrates functions that allow the OS SCSI software stack to discover the iSCSI device, block 916, set and get options/parameters, blocks 903 and 909, to start the device, block 913 and release the device, block 911. Besides the control functions discussed earlier, the iSCSI implementation provides bulk data transfer functions, through queues 912 and 917, to transport the PDUs specified by the iSCSI standard. The iSCSI stack may also include direct data transfer/placement (DDT) or RDMA functions or combination thereof, block 918, which are used by the initiator and target systems to perform substantially zero buffer copy and host intervention-less data transfers including storage and other bulk block data transfers. The SCSI commands and the block data transfers related to these are implemented as command queues, blocks 912 and 917, which get executed on the described processor. The host is interrupted primarily on the command completion. The completed commands are queued for the host to act on at a time convenient to the host. The figure illustrates the iSCSI protocol layer and the driver layer layered on the TCP/IP stack, blocks 907 and 908, which is also implemented off the host processor on the IP processor system described herein.

Fig. 10 illustrates the TCP/IP stack functionality that is implemented in the described IP processor system. These functions provide an interface to the upper layer protocol functions to carry the IP storage traffic as well as other applications that can benefit from direct OS TCP/IP bypass, RDMA or network sockets direct capabilities or combination thereof to utilize the high performance TCP/IP implementation of this processor. The TCP/IP stack provides capabilities to send and receive upper layer data, blocks 1017 and 1031, and command PDUs, establish the transport connections and teardown functions, block 1021, send and receive data transfer functions, checksum functions, block 1019, as well as error handling functions, block 1022, and segmenting and sequencing and windowing operations, block 1023. Certain functions like checksum verification/creation touch every byte of the data transfer whereas some functions that transport the data packets and update the

transmission control block or session data base are invoked for each packet of the data transfer. The session DB, block 1025, is used to maintain various information regarding the active sessions/connections along with the TCP/IP state information. The TCP layer is built on top of IP layer that provides the IP functionality as required by the standard. This layer provides functions to fragment/de-fragment, block 1033, the packets as per the path MTU, providing the route and forwarding information, block 1032, as well as interface to other functions necessary for communicating errors like, for example, ICMP, block 1029. The IP layer interfaces with the Ethernet layer or other media access layer technology to transport the TCP/IP packets onto the network. The lower layer is illustrated as Ethernet in various figures in this description, but could be other technologies like SONET, for instance, to transport the packets over SONET on MANs/WANs. Ethernet may also be used in similar applications, but may be used more so within a LAN and dedicated local SAN environments, for example.

Fig. 11 illustrates the iSCSI data flow. The figure illustrates the receive and transmit path of the data flow. The Host's SCSI command layer working with the iSCSI driver, both depicted in block 1101, would schedule the commands to be processed to the command scheduler, block 1108, in the storage flow controller seen in more detail in Fig. 26. The command scheduler 1108 schedules the new commands for operation in the processor described in more detail in Fig. 17. A new command that is meant for the target device with an existing connection gets en-queued to that existing connection, block 1111. When the connection to the target device does not exist, a new command is en-queued on to the unassigned command queue, block 1102. The session/connection establishment process like that shown in Fig. 47 and blocks 905 and 1006 is then called to connect to the target. Once the connection is established the corresponding command from the queue 1102 gets en-queued to the newly created connection command queue 1111 by the command scheduler 1108 as illustrated in the figure. Once a command reaches a stage of execution, the receive 1107 or transmit 1109 path is activated depending on whether the command is a read or a write transaction. The state of the connection/session which the command is transported is used to record the progress of the command execution in the session database as described subsequently. The buffers associated with the data transfer may be locked till such time as the transfer is completed. If the RDMA mechanism is used to transfer the data between the initiator and the target, appropriate region buffers identifiers, access control keys and related RDMA state data is maintained in memory on board the processor and may also be

maintained in off-chip memory depending on the implementation chosen. As the data transfer, which may be over multiple TCP segments, associated with the command is completed the status of the command execution is passed onto the host SCSI layer which then does the appropriate processing. This may involve releasing the buffers being used for data transfers to the applications, statistics update, and the like. During transfer, the iSCSI PDUs are transmitted by the transmit engines, block 1109, working with the transmit command engines, block 1110, that interpret the PDU and perform appropriate operations like retrieving the application buffers from the host memory using DMA to the storage processor and keeping the storage command flow information in the iSCSI connection database updated with the progress. As used in this patent the term "engine" can be a data processor or a part of a data processor, appropriate for the function or use of the engine. Similarly, the receive engines, block 1107, interpret the received command into new requests, response, errors or other command or data PDUs that need to be acted on appropriately. These receive engines working with the command engines, block 1106, route the read data or received data to the appropriate allocated application buffer through direct data transfer/placement or RDMA control information maintained for the session in the iSCSI session table. On command completion the control to the respective buffers, blocks 1103 and 1112, is released for the application to use. Receive and transmit engines can be the SAN packet processors 1706(a) to 1706(n) of Fig. 17 of this IP processor working with the session information recorded in the session data base entries 1704, which can be viewed as a global memory as viewed from the TCP/IP processor of Fig. 23 or the IP processor of Fig. 24. The same engines can get reused for different packets and commands with the appropriate storage flow context provided by the session database discussed in more detail below with respect to block 1704 and portion of session database in 1708 of Fig. 17. For clarification, the terms IP network application processor, IP Storage processor, IP Storage network application processor and IP processor can be the same entity, depending on the application. An IP network application processor core or an IP storage network application processor core can be the same entity, depending on the application.

Similarly a control command can use the transmit path whereas the received response would use the receive path. Similar engines can exist on the initiator as well as the target. The data flow direction is different depending on whether it is the initiator or the target. However, primarily similar data flow exists on both initiator and target with additional steps at the target. The target needs to perform additional operations to reserve the buffers needed

to get the data of a write command, for instance, or may need to prepare the read data before the data is provided to the initiator. Similar instances would exist in case of an intermediate device, although, in such a device, which may be a switch or an appliance, some level of virtualization or frame filtering or such other operation may be performed that
5 may require termination of the session on one side and originating sessions on the other. This functionality is supported by this architecture but not illustrated explicitly in this figure, inasmuch as it is well within the knowledge of one of ordinary skill in the art.

Fig. 12 through Fig. 15 illustrate certain protocol information regarding transport sessions and how that information may be stored in a database in memory.

10 Fig. 12 illustrates the data structures that are maintained for iSCSI protocol and associated TCP/IP connections. The data belonging to each iSCSI session, block 1201, which is essentially a nexus of initiator and target connections, is carried on the appropriate connection, block 1202. Dependent commands are scheduled on the queues of the same connection to maintain the ordering of the commands, block 1203. However, unrelated
15 commands can be assigned to different transport connection. It is possible to have all the commands be queued to the same connection, if the implementation supports only one connection per session. However, multiple connections per session are feasible to support line trunking between the initiator and the target. For example, in some applications, the initiator and the target will be in communication with each other and will decide through
20 negotiation to accept multiple connections. In others, the initiator and target will communicate through only one session or connection. Fig. 13 and Fig. 14 illustrate the TCP/IP and iSCSI session data base or transmission control block per session and connection. These entries may be carried as separate tables or may be carried together as a composite table as seen subsequently with respect to Figs. 23, 24, 26 and 29 depending
25 on the implementation chosen and the functionality implemented e.g. TCP/IP only, TCP/IP with RDMA, IP Storage only, IP storage with TCP/IP, IP Storage with RDMA and the like. Various engines that perform TCP/IP and storage flow control use all or some of these fields or more fields not shown, to direct the block data transfer over TCP/IP. The appropriate fields are updated as the connection progresses through the multiple states during the
30 course of data transfer. Fig. 15 illustrates one method of storing the transmission control entries in a memory subsystem that consists of an on-chip session cache, blocks 1501 and

1502, and off-chip session memory, blocks 1503, 1504, 1505, 1506 and 1507, that retains the state information necessary for continuous progress of the data transfers.

Fig. 16 illustrates the IP processor architecture at a high level of abstraction. The processor consists of modular and scalable IP network application processor core, block 1603. Its functional blocks provide the functionality for enabling high speed storage and data transport over IP networks. The processor core can include an intelligent flow controller, a programmable classification engine and a storage/network policy engine. Each can be considered an individual processor or any combination of them can be implemented as a single processor. The disclosed processor also includes a security processing block to provide high line rate encryption and decryption functionality for the network packets. This, likewise, can be a single processor, or combined with the others mentioned above. The disclosed processor includes a memory subsystem, including a memory controller interface, which manages the on chip session cache/memory, and a memory controller, block 1602, which manages accesses to the off chip memory which may be SRAM, DRAM, FLASH, ROM, EEPROM, DDR SDRAM, RDRAM, FCRAM, QDR SRAM, or other derivatives of static or dynamic random access memory or a combination thereof. The IP processor includes appropriate system interfaces to allow it to be used in the targeted market segments, providing the right media interfaces, block 1601, for LAN, SAN, WAN and MAN networks, and similar networks, and appropriate host interface, block 1606. The media interface block and the host interface block may be in a multi-port form where some of the ports may serve the redundancy and fail-over functions in the networks and systems in which the disclosed processor is used. The processor also may contain the coprocessor interface block 1605, for extending the capabilities of the main processor for example creating a multi-processor system. The system controller interface of block 1604 allows this processor to interface with an off-the-shelf microcontroller that can act as the system controller for the system in which the disclosed processor may be used. The processor architecture also support a control plane processor on board, that could act as the system controller or session manager. The system controller interface may still be provided to enable the use of an external processor. Such a version of this processor may not include the control processor for die cost reasons. There are various types of the core architecture that can be created, targeting specific system requirements, for example server adapters or storage controllers or switch line cards or other networking systems. The primary differences would be as discussed in the earlier sections of this patent. These processor blocks provide capabilities and performance to

achieve the high performance IP based storage using standard protocols like iSCSI, FCIP, iFCP and the like. The detailed architecture of these blocks will be discussed in the following description.

Fig. 17 illustrates the IP processor architecture in more detail. The architecture provides capabilities to process incoming IP packets from the media access control (MAC) layer, or other appropriate layer, through full TCP/IP termination and deep packet inspection. This block diagram does not show the MAC layer block 1601, or blocks 1602, 1604 or 1605 of Fig. 16. The MAC layer interface blocks to the input queue, block 1701, and output queue, block 1712, of the processor in the media interface, block 1601, shown in Fig. 16. The MAC functionality could be standards based, with the specific type dependent on the network. Ethernet and Packet over SONET are examples of the most widely used interfaces today which may be included on the same silicon or a different version of the processor created with each.

The block diagram in Fig. 17 illustrates input queue and output queue blocks 1701 and 1712 as two separate blocks. The functionality may be provided using a combined block. The input queue block 1701 consists of the logic, control and storage to retrieve the incoming packets from the MAC interface block. Block 1701 queues the packets as they arrive from the interface and creates appropriate markers to identify start of the packet, end of the packet and other attributes like a fragmented packet or a secure packet, and the like, working with the packet scheduler 1702 and the classification engine 1703. The packet scheduler 1702, can retrieve the packets from the input queue controller and passes them for classification to the classification engine. The classification block 1703, is shown to follow the scheduler, however from a logical perspective the classification engine receives the packet from the input queue, classifies the packet and provides the classification tag to the packet, which is then scheduled by the scheduler to the processor array 1706(a) . . . 1706(n). Thus the classification engine can act as a pass-through classification engine, sustaining the flow of the packets through its structure at the full line rate. The classification engine is a programmable engine that classifies the packets received from the network in various categories and tags the packet with the classification result for the scheduler and the other packet processors to use. Classification of the network traffic is a very compute intensive activity which can take up to half of the processor cycles available in a packet processor. This integrated classification engine is programmable to perform Layer

2 through Layer 7 inspection. The fields to be classified are programmed in with expected values for comparison and the action associated with them if there is a match. The classifier collects the classification walk results and can present these as a tag to the packet identifying the classification result as seen subsequently with respect to Fig. 30. This is much like a tree structure and is understood as a "walk." The classified packets are then provided to the scheduler 1702 as the next phase of the processing pipeline:

The packet scheduler block 1702 includes a state controller and sequencer that assign packets to appropriate execution engines on the disclosed processor. The execution engines are the SAN packet processors, block 1706(a) through 1706(n), including the TCP/IP and/or storage engines as well as the storage flow/RDMA controller, block 1708 or host bypass and/or other appropriate processors, depend on the desired implementation. For clarity, the term "/", when used to designate hardware components in this patent, can mean "and/or" as appropriate. For example, the component "storage flow/RDMA controller" can be a storage flow and RDMA controller, a storage flow controller, or an RDMA controller, as appropriate for the implementation. The scheduler 1702 also maintains the packet order through the processor where the state dependency from a packet to a packet on the same connection/session is important for correct processing of the incoming packets. The scheduler maintains various tables to track the progress of the scheduled packets through the processor until packet retirement. The scheduler also receives commands that need to be scheduled to the packet processors on the outgoing commands and packets from the host processor or switch fabric controller or interface.

The TCP/IP and storage engines along with programmable packet processors are together labeled as the SAN Packet Processors 1706(a) through 1706(n) in Fig. 17. These packet processors are engines that are independent programmable entities that serve a specific role. Alternatively, two or more of them can be implemented as a single processor depending on the desired implementation. The TCP/IP engine of Fig. 23 and the storage engines of Fig. 24 are configured in this example as coprocessors to the programmable packet processor engine block 2101 of Fig. 21. This architecture can thus be applied with relative ease to applications other than storage by substituting/removing for the storage engine for reasons of cost, manufacturability, market segment and the like. In a pure networking environment the storage engine could be removed, leaving the packet processor with a dedicated TCP/IP engine and be applied for the networking traffic, which will face the

same processing overhead from TCP/IP software stacks. Alternatively one or more of the engines may be dropped for desired implementation e.g. for processor supporting only IP Storage functions may drop TCP/IP engine and/or packet engine which may be in a separate chip. Hence, multiple variations of the core scalable and modular architecture are possible. The core architecture can thus be leveraged in applications beside the storage over IP applications by substituting the storage engine with other dedicated engines, for example a high performance network security and policy engine, a high performance routing engine, a high performance network management engine, deep packet inspection engine providing string search, an engine for XML, an engine for virtualization, and the like, providing support for an application specific acceleration. The processing capability of this IP processor can be scaled by scaling the number of SAN Packet Processor blocks 1706 (a) through 1706 (n) in the chip to meet the line rate requirements of the network interface. The primary limitation from the scalability would come from the silicon real-estate required and the limits imposed by the silicon process technologies. Fundamentally this architecture is scalable to very high line rates by adding more SAN packet processor blocks thereby increasing the processing capability. Other means of achieving a similar result is to increase the clock frequency of operation of the processor to that feasible within the process technology limits.

Fig. 17 also illustrates the IP session cache/memory and the memory controller block 1704. This cache can be viewed as an internal memory or local session database cache. This block is used to cache and store the TCP/IP session database and also the storage session database for a certain number of active sessions. The number of sessions that can be cached is a direct result of the chosen silicon real-estate and what is economically feasible to manufacture. The sessions that are not on chip, are stored and retrieved to/from off chip memory, viewed as an external memory, using a high performance memory controller block which can be part of block 1704 or otherwise. Various processing elements of this processor share this controller using a high speed internal bus to store and retrieve the session information. The memory controller can also be used to temporarily store packets that may be fragmented or when the host interface or outbound queues are backed-up. The controller may also be used to store statistics information or any other information that may be collected by the disclosed processor or the applications running on the disclosed or host processor.

The processor block diagram of Fig. 17 also illustrates host interface block 1710, host input queue, block 1707 and host output queue, block 1709 as well as the storage flow / RDMA controller, block 1708. These blocks provide the functions that are required to transfer data to and from the host (also called "peer") memory or switch fabric. These blocks also provide
5 features that allow the host based drivers to schedule the commands, retrieve incoming status, retrieve the session database entry, program the disclosed processor, and the like to enable capabilities like sockets direct architecture, full TCP/IP termination, IP storage offload and the like capabilities with or without using RDMA. The host interface controller 1710, seen in greater detail in Fig. 27, provides the configuration registers, DMA engines for direct
10 memory to memory data transfer, the host command block that performs some of the above tasks, along with the host interface transaction controller and the host interrupt controller. The host input and output queues 1707, 1709 provide the queuing for incoming and outgoing packets. The storage flow and RDMA controller block 1708 provides the functionality necessary for the host to queue the commands to the disclosed processor, which then takes
15 these commands and executes them, interrupting the host processor on command termination. The RDMA controller portion of block 1708 provides various capabilities necessary for enabling remote direct memory access. It has tables that include information such as RDMA region, access keys, and virtual address translation functionality. The RDMA engine inside this block performs the data transfer and interprets the received RDMA
20 commands to perform the transaction if the transaction is allowed. The storage flow controller of block 1708 also keeps track of the state of the progress of various commands that have been scheduled as the data transfer happens between the target and the initiator. The storage flow controller schedules the commands for execution and also provides the command completion information to the host drivers. The above can be considered RDMA
25 capability and can be implemented as described or by implementing as individual processors, depending on designer's choice. Also, additional functions can be added to or removed from those described without departing from the spirit or the scope of this patent.

The control plane processor block 1711 of this processor is used to provide relatively slow path functionality for TCP/IP and/or storage protocols which may include error processing
30 with ICMP protocol, name resolution, address resolution protocol, and it may also be programmed to perform session initiation/teardown acting as a session controller/connection manager, login and parameter exchange, and the like. This control plane processor could be off chip to provide the system developer a choice of the control plane processor, or may be

on chip to provide an integrated solution. If the control plane processor is off-chip, then an interface block would be created or integrated herein that would allow this processor to interface with the control plane processor and perform data and command transfers. The internal bus structures and functional block interconnections may be different than illustrated for all the detailed figures for performance, die cost requirements and the like and not depart from the spirit and the scope of this patent.

Capabilities described above for Fig. 17 blocks with more detail below, enable a packet streaming architecture that allows packets to pass through from input to output with minimal latency, with in-stream processing by various processing resources of the disclosed processor.

Fig. 18 illustrates the input queue and controller block shown generally at 1701 of Fig. 17 in more detail. The core functionality of this block is to accept the incoming packets from multiple input ports, Ports 1 to N, in blocks 1801 and 1802(i) to 1802(n), and to queue them using a fixed or programmable priority on the input packet queue, block 1810, from where the packets get de-queued for classifier, scheduler and further packet processing through scheduler I/F blocks 1807-1814. The input queue controller interfaces with each of the input ports (Port 1 through Port N in a multi-port implementation), and queues the packets to the input packet queue 1810. The packet en-queue controller and marker block 1804 may provide fixed priority functions or may be programmable to allow different policies to be applied to different interfaces based on various characteristics like port speed, the network interface of the port, the port priority and others that may be appropriate. Various modes of priority may be programmable like round-robin, weighted round-robin or others. The input packet de-queue controller 1812 de-queues the packets and provides them to the packet scheduler, block 1702 of Fig. 17 via scheduler I/F 1814. The scheduler schedules the packets to the SAN packet processors 1706 (a) – 1706 (n) once the packets have been classified by the classification engine 1703 of Fig. 17. The encrypted packets can be classified as encrypted first and passed on to the security engine 1705 of Fig. 17 by the secure packet interface block 1813 of Fig. 18. for authentication and/or decryption if the implementation includes security processing otherwise the security interfaces may not be present and an external security processor would be used to perform similar functions. The decrypted packets from clear packet interface, block 1811, are then provided to the input queue through block 1812 from which the packet follows the same route as a clear packet.

The fragmented IP packets may be stored on-chip in the fragmented packet store and controller buffers, block 1806, or may be stored in the internal or external memory. When the last fragment arrives, the fragment controller of block 1806, working with the classification engine and the scheduler of Fig. 17, merges these fragments to assemble the complete packet. Once the fragmented packet is combined to form a complete packet, the packet is scheduled into the input packet queue via block 1804 and is then processed by the packet de-queue controller, block 1812, to be passed on to various other processing stages of this processor. The input queue controller of Fig. 18 assigns a packet tag/descriptor to each incoming packet which is managed by the attribute manager of block 1809 which uses the packet descriptor fields like the packet start, size, buffer address, along with any other security information from classification engine, and stored in the packet attributes and tag array of block 1808. The packet tag and attributes are used to control the flow of the packet through the processor by the scheduler and other elements of the processor in an efficient manner through interfaces 1807, 1811, 1813 and 1814

Fig. 19 illustrates the packet scheduler and sequencer 1702 of Fig. 17 in more detail. This block is responsible for scheduling packets and tasks to the execution resources of this processor and thus also acts as a load balancer. The scheduler retrieves the packet headers from the header queue, block 1902, from the input queue controller 1901 to pass them to the classification engine 1703 of Fig. 17 which returns the classification results to the classifier queue, block 1909, that are then used by the rest of the processor engines. The classification engine may be presented primarily with the headers, but if deep packet inspection is also programmed, the classification engine may receive the complete packets which it routes to the scheduler after classification. The scheduler comprises a classification controller/scheduler, block 1908, which manages the execution of the packets through the classification engine. This block 1908 of Fig. 19 provides the commands to the input queue controller, block 1901, in case of fragmented packets or secure packets, to perform the appropriate actions for such packets e.g. schedule an encrypted packet to the security engine of Fig. 17. The scheduler state control and the sequencer, block 1916, receive state information of various transactions/operations active inside the processor and provide instructions for the next set of operations. For instance, the scheduler retrieves the packets from the input packet queue of block 1903, and schedules these packets in the appropriate resource queue depending on the results of the classification received from the classifier or directs the packet to the packet memory, block 1913 or 1704 through 1906, creating a

packet descriptor/tag which may be used to retrieve the packet when appropriate resource needs it to performs its operations at or after scheduling. The state control and sequencer block 1916 instructs/directs the packets with their classification result, block 1914, to be stored in the packet memory, block 1913, from where the packets get retrieved when they are scheduled for operation. The state controller and the sequencer identify the execution resource that should receive the packet for operation and creates a command and assigns this command with the packet tag to the resource queues, blocks 1917 (Control Plane), 1918 (port i-port n), 1919 (bypass) and 1920 (host) of Fig. 19. The priority selector 1921 is a programmable block that retrieves the commands and the packet tag from the respective queues based on the assigned priority and passes this to the packet fetch and command controller, block 1922. This block retrieves the packet from the packet memory store 1913 along with the classification results and schedules the packet transfer to the appropriate resource on the high performance processor command and packet busses such as at 1926 when the resource is ready for operation. The bus interface blocks, like command bus interface controller 1905, of the respective recipients interpret the command and accept the packet and the classification tag for operation. These execution engines inform the scheduler when the packet operation is complete and when the packet is scheduled for its end destination (either the host bus interface, or the output interface or control plane interface, etc.). This allows the scheduler to retire the packet from its state with the help of retirement engine of block 1904 and frees up the resource entry for this session in the resource allocation table, block 1923. The resource allocation table is used by the sequencer to assign the received packets to specific resources, depending on the current state of internal state of these resources, e.g. the session database cache entry buffered in the SAN packet processor engine, the connection ID of the current packet being executed in the resource, and the like. Thus packets that are dependent on an ordered execution get assigned primarily to the same resource, which improves memory traffic and performance by using the current DB state in the session memory in the processor and not have to retrieve new session entries. The sequencer also has interface to the memory controller, block 1906, for queuing of packets that are fragmented packets and/or for the case in which the scheduler queues get backed-up due to a packet processing bottleneck down stream, which may be caused by specific applications that are executed on packets that take more time than that allocated to maintain a full line rate performance, or for the case in which any other downstream systems get full, unable to sustain the line rate.

If the classifier is implemented before the scheduler as discussed above with respect to Fig. 17 where the classification engine receives the packet from the input queue, items 1901, 1902, 1908, 1909 and 1910 would be in the classifier, or may not be needed, depending on the particular design. The appropriate coupling from the classifier to/from the scheduler
5 blocks 1903, 1907, 1914 and 1915 may be created in such a scenario and the classifier coupled directly to the input queue block of Fig. 18.

Fig. 20 illustrates the packet classification engine shown generally at 1703 of Fig. 17. Classification of the packets into their various attributes is a very compute intensive operation. The classifier can be a programmable processor that examines various fields of
10 the received packet to identify the type of the packet, the protocol type e.g. IP, ICMP, TCP, UDP etc, the port addresses, the source and destination fields, etc. The classifier can be used to test a particular field or a set of fields in the header or the payload. The block diagram illustrates a content addressable memory based classifier. However, as discussed earlier this could be a programmable processor as well. The primary differences are the
15 performance and complexity of implementation of the engine. The classifier gets the input packets through the scheduler from the input queues, blocks 2005 and 2004 of Fig. 20. The input buffers 2004 queue the packets/descriptor and/or the packet headers that need to be classified. Then the classification sequencer 2003 fetches the next available packet in the queue and extracts the appropriate packet fields based on the global field descriptor sets,
20 block 2007, which are, or can be, programmed. Then the classifier passes these fields to the content addressable memory (CAM) array, block 2009, to perform the classification. As the fields are passed through the CAM array, the match of these fields identifies next set of fields to be compared and potentially their bit field location. The match in the CAM array results in the action/event tag, which is collected by the result compiler, (where "compiling" is
25 used in the sense of "collecting") block 2014 and also acted on as an action that may require updating the data in the memory array, block 2013, associated with specific CAM condition or rule match. This may include performing an arithmetic logic unit (ALU) operation, block 2017, which can be considered one example of an execution resource) on this field e.g. increment or decrement the condition match and the like. The CAM arrays are programmed
30 with the fields, their expected values and the action on match, including next field to compare, through the database initialization block 2011, accessible for programming through the host or the control plane processor interfaces 1710, 1711. Once the classification reaches a leaf node the classification is complete and the classification tag is generated that

identifies the path traversed that can then be used by other engines of the IP processor avoid performing the same classification tasks. For example a classification tag may include the flow or session ID, protocol type indication e.g. TCP/UDP/ICMP etc., value indicating whether to processes, bypass, drop packet, drop session, and the like, or may also include
5 the specific firmware code routine pointer for the execution resource to start packet processing or may include signature of the classification path traversed or the like. The classification tag fields are chosen based on processor implementation and functionality. The classifier retirement queue, block 2015, holds the packets/descriptors of packets that are classified and classification tag and are waiting to be retrieved by the scheduler. The
10 classification data base can be extended using database extension interface and pipeline control logic block 2006. This allows systems that need extensibility for a larger classification database to be built. The classification engine with the action interpreter, the ALU and range matching block of 2012 also provide capabilities to program storage / network policies / actions that need to be taken if certain policies are met. The
15 policies can be implemented in the form of rule and action tables. The policies get compiled and programmed in the classification engine through the host interface along with the classification tables. The database interface and pipeline control 2006 could be implemented to couple to companion processor to extend the size of the classification/policy engine.

20 Fig. 21 illustrates the SAN Packet Processor shown generally at 1706 (a) through 1706 (n) of Fig. 17. A packet processor can be a specially designed packet processor, or it can be any suitable processor such as an ARM, MIPS, StrongARM, X86, PowerPC, Pentium processor, or any other processor that serves the functions described herein. This is also referred as the packet processor complex in various sections of this patent. This packet
25 processor comprises a packet engine, block 2101, which is generally a RISC machine with target instructions for packet processing or a TCP/IP engine, block 2102 or an IP storage engine, block 2103 or a combination thereof. . These engines can be configured as coprocessors to the packet engine or can be independent engines. Fig. 22 illustrates the packet engine in more detail. The packet engine is a generally RISC machine as indicated
30 above with instruction memory, block 2202, and Data Memory, block 2206, (both of which can be RAM) that are used to hold the packet processing micro routines and the packets and intermediate storage. The instruction memory 2202 which, like all such memory in this patent, can be RAM or other suitable storage, is initialized with the code that is executed

during packet processing. The packet processing code is organized as tight micro routines that fit within the allocated memory. The instruction decoder and the sequencer, block 2204, fetches the instructions from instruction memory 2202, decodes them and sequences them through the execution blocks contained within the ALU, block 2208. This machine can be a
5 simple pipelined engine or a more complex deep pipelined machine that may also be designed to provide a packet oriented instruction set. The DMA engine, block 2205 and the bus controller, block 2201, allow the packet engine to move the data packets from the scheduler of Fig. 19 and the host interface into the data memory 2206 for operation. The DMA engine may hold multiple memory descriptors to store/retrieve packet/data to/from host
10 memory/packet memory. This would enable memory accesses to happen in parallel to packet processor engine operations. The DMA engine 2205 also may be used to move the data packets to and from the TCP and storage engines 2210, 2211. Once the execution of the packet is complete, the extracted data or newly generated packet is transferred to the output interface either towards the media interface or the host interface

15 Fig. 23 illustrates a programmable TCP/IP packet processor engine, seen generally at 2210 of Fig. 22, in more detail. This engine is generally a programmable processor with common RISC instructions along with various TCP/IP oriented instructions and execution engines but could also be a micro-coded or a state machine driven processor with appropriate execution engines described in this patent. The TCP processor includes a checksum block, 2311, for
20 TCP checksum verification and new checksum generation by executing these instructions on the processor. The checksum block extracts the data packet from the packet buffer memory (a Data RAM is one example of such memory), 2309, and performs the checksum generation or verification. The packet look-up interface block, 2310, assists the execution engines and the instruction sequencer, 2305, providing access to various data packet fields or the full data packet. The classification tag interpreter, 2313, is used by the instruction
25 decoder 2304 to direct the program flow based on the results of the classification if such an implementation is chosen. The processor provides specific sequence and windowing operations including segmentation, block 2315, for use in the TCP/IP data sequencing calculations for example, to look-up the next expected sequence number and see if that
30 received is within the agreed upon sliding window, which sliding window is a well known part of the TCP protocol, for the connection to which the packet belongs. This element 2315 may also include a segmentation controller like that show at 2413 of Fig. 24. Alternatively, one of ordinary skill in the art, with the teaching of this patent, can easily implement the

segmentation controllers elsewhere on the TCP/IP processor of this Fig. 23. The processor provides a hash engine, block 2317, which is used to perform hash operations against specific fields of the packet to perform a hash table walk that may be required to get the right session entry for the packet. The processor also includes a register file, block 2316, which
5 extracts various commonly used header fields for TCP processing, along with pointer registers for data source and destination, context register sets, and registers that hold the TCP states along with a general purpose register file. The TCP/IP processor can have multiple contexts for packet execution, so that when a given packet execution stalls for any reason, for example memory access, the other context can be woken up and the processor
10 continue the execution of another packet stream with little efficiency loss. The TCP/IP processor engine also maintains a local session cache, block 2320, which holds most recently used or most frequently used entries, which can be used locally without needing to retrieve them from the global session memory. The local session cache can be considered an internal memory of the TCP/IP processor, which can be a packet processor. Of course,
15 the more entries that will be used that can be stored locally in the internal memory, without retrieving additional ones from the session, or global, memory, the more efficient the processing will be. The packet scheduler of Fig. 19 is informed of the connection IDs that are cached per TCP/IP processor resource, so that it can schedule the packets that belong to the same session to the same packet processor complex. When the packet processor
20 does not hold the session entry for the specific connection, then the TCP session database lookup engine, block 2319, working with the session manager, block 2321, and the hash engine retrieves the corresponding entry from the global session memory through the memory controller interface, block 2323. There are means, such as logic circuitry inside the session manager that allow access of session entries or fields of session entries, that act
25 with the hash engine to generate the session identifier for storing/retrieving the corresponding session entry or its fields to the session database cache. This can be used to update those fields or entries as a result of packet processing. When a new entry is fetched, the entry which it is replacing is stored to the global session memory. The local session caches may follow exclusivity caching principles, so that multiple processor complexes do
30 not cause any race conditions, damaging the state of the session. Other caching protocols like MESI protocol may also be used to achieve similar results. When a session entry is cached in a processor complex, and another processor complex needs that entry, this entry is transferred to the new processor with exclusive access or appropriate caching state based on the algorithm. The session entry may also get written to the global session memory in

certain cases. The TCP/IP processor also includes a TCP state machine, block 2322, which is used to walk through the TCP states for the connection being operated on. This state machine receives the state information stored in the session entry along with the appropriate fields affecting the state from the newly received packet. This allows the state machine to

5 generate the next state if there is a state transition and the information is updated in the session table entry. The TCP/IP processor also includes a frame controller/out of order manager block, 2318, that is used to extract the frame information and perform operations for out of order packet execution. This block could also include an RDMA mechanism such as that shown at 2417 of Fig. 24, but used for non-storage data transfers. One of ordinary

10 skill in the art can also, with the teaching of this patent, implement an RDMA mechanism elsewhere on the TCP/IP processor. This architecture creates an upper layer framing mechanism which may use packet CRC as framing key or other keys that is used by the programmable frame controller to extract the embedded PDUs even when the packets arrive out of order and allow them to be directed to the end buffer destination. This unit interacts

15 with the session database to handle out of order arrival information which is recorded so that once the intermediate segments arrive, the retransmissions are avoided. Once the packet has been processed through the TCP/IP processor, it is delivered for operation to the storage engine, if the packet belongs to a storage data transfer and the specific implementation includes a storage engine, otherwise the packet is passed on to the host

20 processor interface or the storage flow/RDMA controller of block 1708 for processing and for DMA to the end buffer destination. The packet may be transferred to the packet processor block as well for any additional processing on the packet. This may include application and customer specific application code that can be executed on the packet before or after the processing by the TCP/IP processor and the storage processor. Data transfer from the host

25 to the output media interface would also go through the TCP/IP processor to form the appropriate headers to be created around the data and also perform the appropriate data segmentation, working with the frame controller and/or the storage processor as well as to update the session state. This data may be retrieved as a result of host command or received network packet scheduled by the scheduler to the packet processor for operation.

30 The internal bus structures and functional block interconnections may be different than illustrated for performance, die cost requirements and the like. For example, Host Controller Interface 2301, Scheduler Interface 2307 and Memory Controller Interface 2323 may be part of a bus controller that allows transfer of data packets or state information or commands, or a combination thereof, to or from a scheduler or storage flow/RDMA controller or host or

session controller or other resources such as, without limitation, security processor, or media interface units, host interface, scheduler, classification processor, packet buffers or controller processor, or any combination of the foregoing.

Fig. 24 illustrates the IP storage processor engine of Fig. 22 in more detail. The storage engine is a programmable engine with an instruction set that is geared towards IP based storage along with, usually, a normal RISC-like packet processing instruction set. The IP storage processor engine contains block 2411, to perform CRC operations. This block allows CRC generation and verification. The incoming packet with IP storage is transferred from the TCP/IP engine through DMA, blocks 2402 and 2408, into the data memory (a data RAM is an example of such memory), block 2409. When the implementation does not include TCP/IP engine or packet processor engine or a combination thereof, the packet may be received from the scheduler directly for example. The TCP session database information related to the connection can be retrieved from the local session cache as needed or can also be received with the packet from the TCP/IP engine. The storage PDU is provided to the PDU classifier engine, block 2418, which classifies the PDU into the appropriate command, which is then used to invoke the appropriate storage command execution engine, block 2412. The command execution can be accomplished using the RISC, or equivalent, instruction set or using a dedicated hardware engine. The command execution engines perform the command received in the PDU. The received PDU may contain read command data, or R2T for a pending write command or other commands required by the IP storage protocol. These engines retrieve the write data from the host interface or direct the read data to the destination buffer. The storage session database entry is cached, in what can be viewed as a local memory, block 2420, locally for the recent or frequent connections served by the processor. The command execution engines execute the commands and make the storage database entry updates working with the storage state machine, block 2422, and the session manager, block 2421. The connection ID is used to identify the session, and if the session is not present in the cache, then it is retrieved from the global session memory 1704 of Fig. 17 by the storage session look-up engine, block 2419. For data transfer from the initiator to target, the processor uses the segmentation controller, block 2413, to segment the data units into segments as per various network constraints like path MTU and the like. The segmentation controller attempts to ensure that the outgoing PDUs are optimal size for the connection. If the data transfer requested is larger than the maximum effective segment size, then the segmentation controller packs the data into multiple packets and works with

the sequence manager, block 2415, to assign the sequence numbers appropriately. The segmentation controller 2413 may also be implemented within the TCP/IP processor of Fig. 23. That is, the segmentation controller may be part of the sequence/window operations manager 2315 of Fig. 23 when this processor is used for TCP/IP operations and not storage operations. One of ordinary skill in the art can easily suggest alternate embodiments for including the segmentation controller in the TCP/IP processor using the teachings of this patent. The storage processor of Fig. 24 (or the TCP/IP processor of Fig. 23) can also include an RDMA engine that interprets the remote direct memory access instructions received in the PDUs for storage or network data transfers that are implemented using this RDMA mechanism. In Fig. 24, for example, this is RDMA engine 2417. In the TCP/IP processor of Fig. 23 an RDMA engine could be part of the frame controller and out of order manager 2318, or other suitable component. If both ends of the connection agree to the RDMA mode of data transfer, then the RDMA engine is utilized to schedule the data transfers between the target and initiator without substantial host intervention. The RDMA transfer state is maintained in a session database entry. This block creates the RDMA headers to be layered around the data, and is also used to extract these headers from the received packets that are received on RDMA enabled connections. The RDMA engine works with the storage flow/ RDMA controller, 1708, and the host interface controller, 1710, by passing the messages/instructions and performs the large block data transfers without substantial host intervention. The RDMA engine of the storage flow/RDMA controller block, 1708, of the IP processor performs protection checks for the operations requested and also provides conversion from the RDMA region identifiers to the physical or virtual address in the host space. This functionality may also be provided by RDMA engine, block 2417, of the storage engine of the SAN packet processor based on the implementation chosen. The distribution of the RDMA capability between 2417 and 1708 and other similar engines is an implementation choice that one with ordinary skill in the art will be able to do with the teachings of this patent. Outgoing data is packaged into standards based PDU by the PDU creator, block 2425. The PDU formatting may also be accomplished by using the packet processing instructions. The storage engine of Fig. 24 works with the TCP/IP engine of Fig. 23 and the packet processor engine of Fig. 17 to perform the IP storage operations involving data and command transfers in both directions i.e. from the initiator to target and the target to the host and vice versa. That is, the Host controller Interface 2401, 2407 store and retrieve commands or data or a combination thereof to or from the host processor. These interfaces may be directly connected to the host or may be connected through an

intermediate connection. Though shown as two apparatus, interfaces 2401 and 2407 could be implemented as a single apparatus. The flow of data through these blocks would be different based on the direction of the transfer. For instance, when command or data is being sent from the host to the target, the storage processing engines will be invoked first to
5 format the PDU and then this PDU is passed on to the TCP processor to package the PDU in a valid TCP/IP segment. However, a received packet will go through the TCP/IP engine before being scheduled for the storage processor engine. The internal bus structures and functional block interconnections may be different than illustrated for performance, die cost requirements, and the like. For example, and similarly to Fig. 23, Host Controller Interface
10 2401, 2407 and Memory Controller Interface 2423 may be part of a bus controller that allows transfer of data packets or state information or commands, or a combination thereof, to or from a scheduler or host or storage flow/RDMA controller or session controller or other resources such as, without limitation, security processor, or media interface units, host interface, scheduler, classification processor, packet buffers or controller processor, or any
15 combination of the foregoing.

In applications in which storage is done on a chip not including the TCP/IP processor of Fig. 23 by, as one example, an IP Storage processor such as an iSCSI processor of Fig. 24, the TCP/IP Interface 2406 would function as an interface to a scheduler for scheduling IP storage packet processing by the IP Storage processor. Similar variations are well within the
20 knowledge of one of ordinary skill in the art, viewing the disclosure of this patent.

Fig. 25 illustrates the output queue controller block 1712 of Fig. 17 in more detail. This block receives the packets that need to be sent on to the network media independent interface 1601 of Fig. 16. The packets may be tagged to indicate if they need to be encrypted before being sent out. The controller queues the packets that need to be secured to the security
25 engine through the queue 2511 and security engine interface 2510. The encrypted packets are received from the security engine and are queued in block 2509, to be sent to their destination. The output queue controller may assign packets onto their respective quality of service (QOS) queues, if such a mechanism is supported. The programmable packet priority selector, block 2504, selects the next packet to be sent and schedules the packet for
30 the appropriate port, Port1 . . . PortN. The media controller block 1601 associated with the port accepts the packets and sends them to their destination.

Fig. 26 illustrates the storage flow controller /RDMA controller block, shown generally at 1708 of Fig. 17, in more detail. The storage flow and RDMA controller block provides the functionality necessary for the host to queue the commands (storage or RDMA or sockets direct or a combination thereof) to this processor, which then takes these commands and executes them, interrupting the host processor primarily on command termination. The command queues, new and active, blocks 2611 and 2610, and completion queue, block 2612, can be partially on chip and partially in a host memory region or memory associated with the IP processor, from which the commands are fetched or the completion status deposited. The RDMA engine, block 2602, provides various capabilities necessary for enabling remote direct memory access. It has tables, like RDMA look-up table 2608, that include information like RDMA region and the access keys, and virtual address translation functionality. The RDMA engine inside this block 2602 performs the data transfer and interprets the received RDMA commands to perform the transaction if allowed. The storage flow controller also keeps track of the state of the progress of various commands that have been scheduled as the data transfer happens between the target and the initiator. The storage flow controller schedules the commands for execution and also provides the command completion information to the host drivers. The storage flow controller provides command queues where new requests from the host are deposited, as well as active commands are held in the active commands queue. The command scheduler of block 2601, assigns new commands, that are received which are for targets for which no connections exist, to the scheduler for initiating a new connection. The scheduler 1702, uses the control plane processor shown generally at 1711 of Fig. 17 to do the connection establishment at which point the connection entry is moved to the session cache, shown generally in Fig. 15 and 1704 in Fig. 17, and the state controller in the storage flow controller block 2601 moves the new command to active commands and associates the command to the appropriate connection. The active commands, in block 2610, are retrieved and sent to the scheduler, block 1702 for operation by the packet processors. The update to the command status is provided back to the flow controller which then stores it in the command state tables, blocks 2607 and accessed through block 2603. The sequencer of 2601 applies a programmable priority for command scheduling and thus selects the next command to be scheduled from the active commands and new commands. The flow controller also includes a new requests queue for incoming commands, block 2613. The new requests are transferred to the active command queue once the appropriate processing and buffer reservations are done on the host by the host driver. As the commands are being scheduled for execution, the state

controller 2601 initiates data pre-fetch by host data pre-fetch manager, block 2617, from the host memory using the DMA engine of the host interface block 2707, hence keeping the data ready to be provided to the packet processor complex when the command is being executed. The output queue controller, block 2616, enables the data transfer, working with
5 the host controller interface, block 2614. The storage flow/RDMA controller maintains a target-initiator table, block 2609, that associates the target/initiators that have been resolved and connections established for fast look-ups and for associating commands to active connections. The command sequencer may also work with the RDMA engine 2602, if the commands being executed are RDMA commands or if the storage transfers were negotiated
10 to be done through the RDMA mechanism at the connection initiation. The RDMA engine 2602, as discussed above, provides functionality to accept multiple RDMA regions, access control keys and the virtual address translation pointers. The host application (which may be a user application or an OS kernel function, storage or non-storage such as downloading web pages, video files, or the like) registers a memory region that it wishes to use in RDMA
15 transactions with the disclosed processor through the services provided by the associated host driver. Once this is done, the host application communicates this information to its peer on a remote end. Now, the remote machine or the host can execute RDMA commands, which are served by the RDMA blocks on both ends without requiring substantial host intervention. The RDMA transfers may include operations like read from a region, a certain
20 number of bytes with a specific offset or a write with similar attributes. The RDMA mechanism may also include send functionality which would be useful in creating communication pipes between two end nodes. These features are useful in clustering applications where large amounts of data transfer is required between buffers of two applications running on servers in a cluster, or more likely, on servers in two different
25 clusters of servers, or such other clustered systems. The storage data transfer may also be accomplished using the RDMA mechanism, since it allows large blocks of data transfers without substantial host intervention. The hosts on both ends get initially involved to agree on doing the RDMA transfers and allocating memory regions and permissions through access control keys that get shared. Then the data transfer between the two nodes can
30 continue without host processor intervention, as long as the available buffer space and buffer transfer credits are maintained by the two end nodes. The storage data transfer protocols would run on top of RDMA, by agreeing to use RDMA protocol and enabling it on both ends. The storage flow controller and RDMA controller of Fig. 26 can then perform the storage command execution and the data transfer using RDMA commands. As the

- expected data transfers are completed the storage command completion status is communicated to the host using the completion queue 2612. The incoming data packets arriving from the network are processed by the packet processor complex of Fig. 17 and then the PDU is extracted and presented to the flow controller OF FIG. 26 in case of
- 5 storage/RDMA data packets. These are then assigned to the incoming queue block 2604, and transferred to the end destination buffers by looking up the memory descriptors of the receiving buffers and then performing the DMA using the DMA engine inside the host interface block 2707. The RDMA commands may also go through protection key look-up and address translation as per the RDMA initialization.
- 10 The foregoing may also be considered a part of an RDMA capability or an RDMA mechanism or an RDMA function.

- Fig. 27 illustrates host interface controller 1710 of Fig. 17 in more detail. The host interface block includes a host bus interface controller, block 2709, which provides the physical interface to the host bus. The host interface block may be implemented as a fabric interface
- 15 or media independent interface when embodied in a switch or a gateway or similar configuration depending on the system architecture and may provide virtual output queuing and/or other quality of service features. The transaction controller portion of block 2708, executes various bus transactions and maintains their status and takes requested transactions to completion. The host command unit, block 2710, includes host bus
- 20 configuration registers and one or more command interpreters to execute the commands being delivered by the host. The host driver provides these commands to this processor over Host Output Queue Interface 2703. The commands serve various functions like setting up configuration registers, scheduling DMA transfers, setting up DMA regions and permissions if needed, setup session entries, retrieve session database, configure RDMA
- 25 engines and the like. The storage and other commands may also be transferred using this interface for execution by the IP processor.

- Fig. 28 illustrates the security engine 1705 of Fig. 17 in more detail. The security engine illustrated provides authentication and encryption and decryption services like those required by standards like IPSEC for example. The services offered by the security engine may
- 30 include multiple authentication and security algorithms. The security engine may be on-board the processor or may be part of a separate silicon chip as indicated earlier. An external security engine providing IP security services would be situated in a similar position

in the data flow, as one of the first stages of packet processing for incoming packets and as one of the last stages for the outgoing packet. The security engine illustrated provides advanced encryption standard (AES) based encryption and decryption services, which are very hardware performance efficient algorithms adopted as security standards. This block could also provide other security capabilities like DES, 3DES, as an example. The supported algorithms and features for security and authentication are driven from the silicon cost and development cost. The algorithms chosen would also be those required by the IP storage standards. The authentication engine, block 2803, is illustrated to include the SHA-1 algorithm as one example of useable algorithms. This block provides message digest and authentication capabilities as specified in the IP security standards. The data flows through these blocks when security and message authentication services are required. The clear packets on their way out to the target are encrypted and are then authenticated if required using the appropriate engines. The secure packets received go through the same steps in reverse order. The secure packet is authenticated and then decrypted using the engines 2803, 2804 of this block. The security engine also maintains the security associations in a security context memory, block 2809, that are established for the connections. The security associations (may include secure session index, security keys, algorithms used, current state of session and the like) are used to perform the message authentication and the encryption/decryption services. It is possible to use the message authentication service and the encryption/decryption services independent of each other.

Fig. 29 illustrates the session cache and memory controller complex seen generally at 1704 of Fig. 17 in more detail. The memory complex includes a cache/memory architecture for the TCP/IP session database called session/global session memory or session cache in this patent, implemented as a cache or memory or a combination thereof. The session cache look-up engine, block 2904, provides the functionality to look-up a specific session cache entry. This look-up block creates a hash index out of the fields provided or is able to accept a hash key and looks-up the session cache entry. If there is no tag match in the cache array with the hash index, the look-up block uses this key to find the session entry from the external memory and replaces the current session cache entry with that session entry. It provides the session entry fields to the requesting packet processor complex. The cache entries that are present in the local processor complex cache are marked shared in the global cache. Thus when any processor requests this cache entry, it is transferred to the global cache and the requesting processor and marked as such in the global cache. The

session memory controller is also responsible to move the evicted local session cache entries into the global cache inside this block. Thus only the latest session state is available at any time to any requesters for the session entry. If the session cache is full, a new entry may cause the least recently used entry to be evicted to the external memory. The session
5 memory may be single way or multi-way cache or a hash indexed memory or a combination thereof, depending on the silicon real estate available in a given process technology. The use of a cache for storing the session database entry is unique, in that in networking applications for network switches or routers, generally there is not much locality of reference properties available between packets, and hence use of cache may not provide much
10 performance improvement due to cache misses. However, the storage transactions are longer duration transactions between the two end systems and may exchange large amounts of data. In this scenario or cases where a large amount of data transfer occurs between two nodes, like in clustering or media servers or the like a cache based session memory architecture will achieve significant performance benefit from reducing the
15 enormous data transfers from the off chip memories. The size of the session cache is a function of the available silicon die area and can have an impact on performance based on the trade-off. The memory controller block also provides services to other blocks that need to store packets, packet fragments or any other operating data in memory. The memory interface provides single or multiple external memory controllers, block 2901, depending on
20 the expected data bandwidth that needs to be supported. This can be a double data rate controller or controller for DRAM or SRAM or RDRAM or other dynamic or static RAM or combination thereof. The figure illustrates multi-controllers however the number is variable depending on the necessary bandwidth and the costs. The memory complex may also provide timer functionality for use in retransmission time out for sessions that queue
25 themselves on the retransmission queues maintained by the session database memory block.

Fig. 30 illustrates the data structures details for the classification engine. This is one way of organizing the data structures for the classification engine. The classification database is illustrated as a tree structure, block 3001, with nodes, block 3003, in the tree and the
30 actions, block 3008, associated with those nodes allow the classification engine to walk down the tree making comparisons for the specific node values. The node values and the fields they represent are programmable. The action field is extracted when a field matches a specific node value. The action item defines the next step, which may include extracting and

comparing a new field, performing other operations like ALU operations on specific data fields associated with this node-value pair, or may indicate a terminal node, at which point the classification of the specific packet is complete. This data structure is used by the classification engine to classify the packets that it receives from the packet scheduler. The action items that are retrieved with the value matches, while iterating different fields of the packet, are used by the results compiler to create a classification tag, which is attached to the packet, generally before the packet headers. The classification tag is then used as a reference by the rest of the processor to decide on the actions that need to be taken based on the classification results. The classifier with its programmable characteristics allows the classification tree structure to be changed in-system and allow the processor to be used in systems that have different classification needs. The classification engine also allows creation of storage /network policies that can be programmed as part of the classification tree-node-value-action structures and provide a very powerful capability in the IP based storage systems. The policies would enhance the management of the systems that use this processor and allow enforcement capabilities when certain policies or rules are met or violated. The classification engine allows expansion of the classification database through external components, when that is required by the specific system constraints. The number of trees and nodes are decided based on the silicon area and performance tradeoffs. The data structure elements are maintained in various blocks of the classification engine and are used by the classification sequencer to direct the packet classification through the structures. The classification data structures may require more or less fields than those indicated depending on the target solution. Thus the core functionality of classification may be achieved with fewer components and structures without departing from the basic architecture. The classification process walks through the trees and the nodes as programmed. A specific node action may cause a new tree to be used for the remaining fields for classification. Thus, the classification process starts at the tree root and progress through the nodes until it reaches the leaf node.

Fig. 31 illustrates a read operation between an initiator and target. The initiator sends a READ command request, block 3101, to the target to start the transaction. This is an application layer request which is mapped to specific SCSI protocol command which is then transported as an READ protocol data unit, block 3102, in an IP based storage network. The target prepares the data that is requested, block 3103 and provides read response PDUs, block 3105, segmented to meet the maximum transfer unit limits. The initiator then

retrieves the data, block 3016, from the IP packets and is then stored in the read buffers allocated for this operation. Once all the data has been transferred the target responds with command completion and sense status, block 3107. The initiator then retires the command once the full transfer is complete, block 3109. If there were any errors at the target and the
5 command is being aborted for any reason, then a recovery procedure may be initiated separately by the initiator. This transaction is a standard SCSI READ transaction with the data transport over IP based storage protocol like iSCSI as the PDUs of that protocol.

Fig. 32 illustrates the data flow inside the IP processor of this invention for one of the received READ PDUs of the transaction illustrated in Fig. 31. The internal data flow is
10 shown for the read data PDU received by the IP processor on the initiator end. This figure illustrates various stage of operation that a packet goes through. The stages can be considered as pipeline stages through which the packets traverse. The number of pipe stages traversed depends on the type of the packet received. The figure illustrates the pipe stages for a packet received on an established connection. The packet traverses through
15 the following major pipe stages:

1. Receive Pipe Stage of block 3201, with major steps illustrated in block 3207:
Packet is received by the media access controller. The packet is detected, the preamble/
trailers removed and a packet extracted with the layer2 header and the payload. This is the
stage where the Layer2 validation occurs for the intended recipient as well as any error
20 detection. There may be quality of service checks applied as per the policies established.
Once the packet validation is clear the packet is queued to the input queue.

2. Security Pipe Stage of block 3202, with major steps illustrated in block 3208.
The packet is moved from the input queue to the classification engine, where a quick
determination for security processing is made and if the packet needs to go through security
25 processing, it enters the security pipe stage. If the packet is received in clear text and does not need authentication, then the security pipe stage is skipped. The security pipe stage may also be omitted if the security engine is not integrated with the IP processor. The packet goes through various stages of security engine where first the security association for this connection is retrieved from memory, and the packet is authenticated using the
30 message authentication algorithm selected. The packet is then decrypted using the security keys that have been established for the session. Once the packet is in clear text, it is queued back to the input queue controller.

3. Classification Pipe Stage of block 3203, with major steps illustrated in block 3209. The scheduler retrieves the clear packet from the input queue and schedules the packet for classification. The classification engine performs various tasks like extracting the relevant fields from the packet for layer 3 and higher layer classification, identifies TCP/IP/ storage protocols and the like and creates those classification tags and may also take actions like rejecting the packet or tagging the packet for bypass depending on the policies programmed in the classification engine. The classification engine may also tag the packet with the session or the flow to which it belongs along with marking the packet header and payload for ease of extraction. Some of the tasks listed may be or may not be performed and other tasks may be performed depending on the programming of the classification engine. As the classification is done, the classification tag is added to the packet and packet is queued for the scheduler to process.

4. Schedule Pipe Stage of block 3204, with major steps illustrated in block 3210. The classified packet is retrieved from the classification engine queue and stored in the scheduler for it to be processed. The scheduler performs the hash of the source and destination fields from the packet header to identify the flow to which the packet belongs, if not done by the classifier. Once the flow identification is done the packet is assigned to an execution resource queue based on the flow dependency. As the resource becomes available to accept a new packet, the next packet in the queue is assigned for execution to that resource.

5. Execution Pipe Stage of block 3205, with major steps illustrated in block 3211. The packet enters the execution pipe stage when the resource to execute this packet becomes available. The packet is transferred to the packet processor complex that is supposed to execute the packet. The processor looks at the classification tag attached to the packet to decide the processing steps required for the packet. If this is an IP based storage packet, then the session database entry for this session is retrieved. The database access may not be required if the local session cache already holds the session entry. If the packet assignment was done based on the flow, then the session entry may not need to be retrieved from the global session memory. The packet processor then starts the TCP engine/ the storage engines to perform their operations. The TCP engine performs various TCP checks including checksum, sequence number checks, framing checks with necessary CRC operations, and TCP state update. Then the storage PDU is extracted and assigned to

the storage engine for execution. The storage engine interprets the command in the PDU and in this particular case identifies it to be a read response for an active session. It then verifies the payload integrity and the sequence integrity and then updates the storage flow state in the session database entry. The memory descriptor of the destination buffer is also
5 retrieved from the session data base entry and the extracted PDU payload is queued to the storage flow/RDMA controller and the host interface block for them to DMA the data to the final buffer destination. The data may be delivered to the flow controller with the memory descriptor and the command/operation to perform. In this case deposit the data for this active read command. The storage flow controller updates its active command database.
10 The execution engine indicates to the scheduler the packet has been retired and the packet processor complex is ready to receive its next command.

6. DMA Pipe Stage of block 3206, with major steps illustrated in block 3212. Once the storage flow controller makes the appropriate verification of the Memory descriptor, the command and the flow state, it passes the data block to the host DMA engine for transfer
15 to the host memory. The DMA engine may perform priority based queuing, if such QOS mechanism is programmed or implemented. The data is transferred to the host memory location through DMA. If this is the last operation of the command, then the command execution completion is indicated to the host driver. If this is the last operation for a command and the command has been queued to the completion queue, the resources
20 allocated for the command are released to accept new command. The command statistics may be collected and transferred with the completion status as may be required for performance analysis, policy management or other network management or statistical purposes.

Fig. 33 illustrates write command operation between an initiator and a target. The Initiator
25 sends a WRITE command, block 3301, to the target to start the transaction. This command is transported as a WRITE PDU, block 3302, on the IP storage network. The receiver queues the received command in the new request queue. Once the old commands in operation are completed, block 3304, the receiver allocates the resources to accept the WRITE data corresponding to the command, block 3305. At this stage the receiver issues a
30 ready to transfer (R2T) PDU, block 3306, to the initiator, with indication of the amount of data it is willing to receive and from which locations. The initiator interprets the fields of the R2T requests and sends the data packets, block 3307, to the receiver as per the received R2T.

This sequence of exchange between the initiator and target continues until the command is terminated. A successful command completion or an error condition is communicated to the initiator by the target as a response PDU, which then terminates the command. The initiator may be required to start a recovery process in case of an error. This is not shown in the exchange of the Fig. 33.

Fig. 34 illustrates the data flow inside the IP processor of this invention for one of the R2T PDUs and the following write data of the write transaction illustrated in Fig. 33. The initiator receives the R2T packet through its network media interface. The packet passes through all the stages, blocks 3401, 3402, 3403, and 3404 with detailed major steps in corresponding blocks 3415, 3416, 3409 and 3410, similar to the READ PDU in Fig. 32 including Receive, Security, Classification, Schedule, and Execution. Security processing is not illustrated in this figure. Following these stages the R2T triggers the write data fetch using the DMA stage shown in Fig. 34, blocks 3405 and 3411. The write data is then segmented and put in TCP/IP packets through the execution stage, blocks 3406 and 3412. The TCP and storage session DB entries are updated for the WRITE command with the data transferred in response to the R2T. The packet is then queued to the output queue controller. Depending on the security agreement for the connection, the packet may enter the security pipe stage, block 3407 and 3413. Once the packet has been encrypted and message authentication codes generated, the packet is queued to the network media interface for the transmission to the destination. During this stage, block 3408 and 3414 the packet is encapsulated in the Layer 2 headers, if not already done so by the packet processor and is transmitted. The steps followed in each stage of the pipeline are similar to that of the READ PDU pipe stages above, with additional stages for the write data packet stage, which is illustrated in this figure. The specific operations performed in each stage depend on the type of the command, the state of the session, the command state and various other configurations for policies that may be setup.

Fig. 35 illustrates the READ data transfer using RDMA mechanism between and initiator and target. The initiator and target register the RDMA buffers before initiating the RDMA data transfer, blocks 3501, 3502, and 3503. The initiator issues a READ command, block 3510, with the RDMA buffer as the expected recipient. This command is transported to the target, block 3511. The target prepares the data to be read, block 3504, and then performs the RDMA write operations, block 3505 to directly deposit the read data into the RDMA buffers

at the initiator without the host intervention. The operation completion is indicated using the command completion response.

Fig. 36 illustrates the internal architecture data flow for the RDMA Write packet implementing the READ command flow. The RDMA write packet also follows the same pipe stages as any other valid data packet that is received on the network interface. This packet goes through Layer 2 processing in the receive pipe stage, blocks 3601 and 3607, from where it is queued for scheduler to detect the need for security processing. If the packet needs to be decrypted or authenticated, it enters the security pipe stage, blocks 3602 and 3608. The decrypted packet is then scheduled to the classification engine for it to perform the classification tasks that have been programmed, blocks 3603 and 3609. Once classification is completed, the tagged packet enters the schedule pipe stage, blocks 3604 and 3610, where the scheduler assigns this packet to a resource specific queue dependent on flow based scheduling. When the intended resource is ready to execute this packet, it is transferred to that packet processor complex, blocks 3605 and 3611, where all the TCP/IP verification, checks, and state updates are made and the PDU is extracted. Then the storage engine identifies the PDU as belonging to a storage flow for storage PDUs implemented using RDMA and interprets the RDMA command. In this case it is RDMA write to a specific RDMA buffer. This data is extracted and passed on to the storage flow/RDMA controller block which performs the RDMA region translation and protection checks and the packet is queued for DMA through the host interface, blocks 3606 and 3612. Once the packet has completed operation through the packet processor complex, the scheduler is informed and the packet is retired from the states carried in the scheduler. Once in the DMA stage, the RDMA data transfer is completed and if this is the last data transfer that completes the storage command execution, that command is retired and assigned to the command completion queue.

Fig. 37 illustrates the storage write command execution using RDMA Read operations. The initiator and target first register their RDMA buffers with their RDMA controllers and then also advertise the buffers to their peer. Then the initiator issues a write command, block 3701, to the target, where it is transported using the IP storage PDU. The recipient executes the write command, by first allocating the RDMA buffer to receive the write and then requesting an RDMA read to the initiator, blocks 3705, and 3706. The data to be written from the initiator is then provided as an RDMA read response packet, blocks 3707 and 3708. The receiver deposits the packet directly to the RDMA buffer without any host interaction. If the

read request was for data larger than the segment size, then multiple READ response PDUs would be sent by the initiator in response to the READ request. Once the data transfer is complete the completion status is transported to the initiator and the command completion is indicated to the host.

- 5 Fig. 38 illustrates the data flow of an RDMA Read request and the resulting write data transfer for one section of the flow transaction illustrated in Fig. 37. The data flow is very similar to the write data flow illustrated in Fig. 34. The RDMA read request packet flows through various processing pipe stages including: receive, classify, schedule, and execution, blocks 3801, 3802, 3803, 3804, 3815, 3816, 3809 and 3810. Once this request is executed, it generates the RDMA read response packet. The RDMA response is generated by first
10 doing the DMA, blocks 3805 and 3811, of the requested data from the system memory, and then creating segments and packets through the execution stage, blocks 3806 and 3812. The appropriate session database entries are updated and the data packets go to the security stage, if necessary, blocks 3807 and 3813. The secure or clear packets are then
15 queued to the transmit stage, block 3808 and 3814, which performs the appropriate layer 2 updates and transmits the packet to the target.

Fig. 39 illustrates an initiator command flow for the storage commands initiated from the initiator in more details. As illustrated following are some of the major steps that a command follows:

- 20 1. Host driver queues the command in processor command queue in the storage flow/RDMA controller;
2. Host is informed if the command is successfully scheduled for operation and to reserve the resources;
- 25 3. The storage flow/RDMA controller schedules the command for operation to the packet scheduler, if the connection to the target is established. Otherwise the controller initiates the target session initiation and once session is established the command is scheduled to the packet scheduler;
4. The scheduler assigns the command to one of the SAN packet processors that is ready to accept this command;

5. The processor complex sends a request to the session controller for the session entry;
6. The session entry is provided to the packet processor complex;
7. The packet processor forms a packet to carry the command as a PDU and is scheduled to the output queue; and
8. The command PDU is given to the network media interface, which sends it to the target.

This is the high level flow primarily followed by most commands from the initiator to the target when the connection has been established between an initiator and a target.

- 10 Fig. 40 illustrates read packet data flow in more detail. Here the read command is initially send using a flow similar to that illustrated in Fig. 39 from the initiator to the target. The target sends the read response PDU to the initiator which follows the flow illustrated in Fig. 40. As illustrated the read data packet passes through following major steps:

1. Input packet is received from the network media interface block;
- 15 2. Packet scheduler retrieves the packet from the input queue;
3. Packet is scheduled for classification;
4. Classified packet returns from the classifier with a classification tag;
5. Based on the classification and flow based resource allocation, the packet is assigned to a packet processor complex which operates on the packet;
- 20 6. Packet processor complex looks-up session entry in the session cache (if not present locally);
7. Session cache entry is returned to the packet processor complex;
8. Packet processor complex performs the TCP/IP operations / IP storage operations and extracts the read data in the payload. The read data with appropriate

destination tags like MDL(memory descriptor list) is provided to the host interface output controller; and

9. The host DMA engine transfers the read data to the system buffer memory.

5 Some of these steps are provided in more details in Fig. 32, where a secure packet flow is represented, where as the Fig. 40 represents a clear text read packet flow. This flow and other flows illustrated in this patent are applicable to storage and non-storage data transfers by using appropriate resources of the disclosed processor, that a person with ordinary skill in the art will be able to do with the teachings of this patent.

10 Fig. 41 illustrates the write data flow in more details. The write command follows the flow similar to that in Fig. 39. The initiator sends the write command to the target. The target responds to the initiator with a ready to transfer (R2T) PDU which indicates to the initiator that the target is ready to receive the specified amount of data. The initiator then sends the requested data to the target. Fig. 41 illustrates the R2T followed by the requested write data packet from the initiator to the target. The major steps followed in this flow are as follows:

- 15 1. Input packet is received from the network media interface block;
2. Packet scheduler retrieves the packet from the input queue;
3. Packet is scheduled for classification;
4. Classified packet returns from the classifier with a classification tag;
 - a. Depending on the classification and flow based resource allocation,
20 the packet is assigned to a packet processor complex which operates on the packet;
5. Packet processor complex looks-up session entry in the session cache (if not present locally);
6. Session cache entry is returned to the packet processor complex;
7. The packet processor determines the R2T PDU and requests the write data
25 with a request to the storage flow/RDMA Controller;

8. The flow controller starts the DMA to the host interface;
 9. Host interface performs the DMA and returns the data to the host input queue;
 10. The packet processor complex receives the data from the host input queue;
 - 5 11. The packet processor complex forms a valid PDU and packet around the data, updates the appropriate session entry and transfers the packet to the output queue; and
 12. The packet is transferred to the output network media interface block which transmits the data packet to the destination.
- 10 The flow in Fig. 41 illustrates clear text data transfer. If the data transfer needs to be secure, the flow is similar to that illustrated in Fig. 43, where the output data packet is routed through the secure packet as illustrated by arrows labeled 11a and 11b. The input R2T packet, if secure would also be routed through the security engine (this is not illustrated in the figure).
- Fig. 42 illustrates the read packet flow when the packet is in cipher text or is secure. This flow is illustrated in more details in Fig. 32 with its associated description earlier. The primary difference between the secure read flow and the clear read flow is that the packet is initially classified as secure packet by the classifier, and hence is routed to the security engine. These steps are illustrated by arrows labeled 2a, 2b, and 2c. The security engine decrypts the packet and performs the message authentication, and transfers the clear packet to the input queue for further processing as illustrated by arrow labeled 2d. The clear packet is then retrieved by the scheduler and provided to the classification engine as illustrated by arrows labeled 2e and 3 in Fig. 42. The rest of the steps and operations are the same as that in Fig. 40, described above.
- 25 Fig. 44 illustrates the RDMA buffer advertisement flow. This flow is illustrated to be very similar to any other storage command flow as illustrated in the Fig. 39. The detailed actions taken in the major steps are different depending on the command. For RDMA buffer advertisement and registration, the RDMA region id is created and recorded along with the address translation mechanism for this region is recorded. The RDMA registration also

includes the protection key for the access control and may include other fields necessary for RDMA transfer. The steps to create the packet for the command are similar to those of Fig. 39.

5 Fig. 45 illustrates the RDMA write flow in more details. The RDMA writes appear like normal read PDUs to the initiator receiving the RDMA write. The RDMA write packet follows the same major flow steps as a read PDU illustrated in Fig. 40. The RDMA transfer involves the RDMA address translation and region access control key checks, and updating the RDMA database entry, beside the other session entries. The major flow steps are the same as the regular Read response PDU.

10 Fig. 46 illustrates the RDMA Read data flow in more details. This diagram illustrates the RDMA read request being received by the initiator from the target and the RDMA Read data being written out from the initiator to the target. This flow is very similar to the R2T response followed by the storage write command. In this flow the storage write command is accomplished using RDMA Read. The major steps that the packet follows are primarily the
15 same as the R2T/write data flow illustrated in Fig. 41.

Fig. 47 illustrates the major steps of session creation flow. This figure illustrates the use of the control plane processor for this slow path operation required at the session initiation between an initiator and a target. This functionality is possible to implement through the packet processor complex. However, it is illustrated here as being implemented using the
20 control plane processor. Both approaches are acceptable. Following are the major steps during session creation:

1. The command is scheduled by the host driver;
2. The host driver is informed that the command is scheduled and any control information required by the host is passed;
- 25 3. The storage flow/RDMA controller detects a request to send the command to a target for which a session is not existing, and hence it passes the request to the control plane processor to establish the transport session;
4. Control plane processor sends a TCP SYN packet to the output queue;

5. The SYN packet is transmitted to the network media interface from which is transmitted to the destination;
6. The destination, after receiving the SYN packet, responds with the SYN-ACK response, which packet is queued in the input queue on receipt from the network media interface;
7. The packet is retrieved by the packet scheduler;
8. The packet is passed to the classification engine;
9. The tagged classified packet is returned to the scheduler;
10. The scheduler, based on the classification, forwards this packet to control plane processor;
11. The processor then responds with an ACK packet to the output queue;
12. The packet is then transmitted to the end destination thus finishing the session establishment handshake; and
13. Once the session is established, this state is provided to the storage flow controller. The session entry is thus created which is then passed to the session memory controller (this part not illustrated in the figure).

Prior to getting the session in the established state as in step 13, the control plane processor may be required to perform a full login phase of the storage protocol, exchanging parameters and recording them for the specific connection if this is a storage data transfer connection. Once the login is authenticated and parameter exchange complete, does the session enter the session establishment state shown in step 13 above.

Fig. 48 illustrates major steps in the session tear down flow. The steps in this flow are very similar to those in Fig. 47. Primary difference between the two flows is that, instead of the SYN, SYN-ACK and ACK packets for session creation, FIN, FIN-ACK and ACK packets are transferred between the initiator and the target. The major steps are otherwise very similar. Another major difference here is that the appropriate session entry is not created but

removed from the session cache and the session memory. The operating statistics of the connection are recorded and may be provided to the host driver, although this is not illustrated in the figure.

Fig. 49 illustrates the session creation and session teardown steps from a target perspective.

5 Following are the steps followed for the session creation:

1. The SYN request from the initiator is received on the network media interface;
2. The scheduler retrieves the SYN packet from the input queue;
3. The scheduler sends this packet for classification to the classification engine;
4. The classification engine returns the classified packet with appropriate tags;
- 10 5. The scheduler, based on the classification as a SYN packet, transfers this packet to the control plane processor;
6. Control plane processor responds with a SYN-ACK acknowledgement packet. It also requests the host to allocate appropriate buffer space for unsolicited data transfers from the initiator (this part is not illustrated);
- 15 7. The SYN-ACK packet is sent to the initiator;
8. The initiator then acknowledges the SYN-ACK packet with an ACK packet, completing the three-way handshake. This packet is received at the network media interface and queued to the input queue after layer 2 processing;
9. The scheduler retrieves this packet;
- 20 10. The packet is sent to the classifier;
11. Classified packet is returned to the scheduler and is scheduled to be provided to the control processor to complete the three way handshake;
12. The controller gets the ACK packet;

13. The control plane processor now has the connection in an established state and it passes the to the storage flow controller which creates the entry in the session cache; and
14. The host driver is informed of the completed session creation.
- 5 The session establishment may also involve the login phase, which is not illustrated in the Fig. 49. However, the login phase and the parameter exchange occur before the session enters the fully configured and established state. These data transfers and handshake may primarily be done by the control processor. Once these steps are taken the remaining steps in the flow above may be executed.
- 10 Figs. 50 and 51 illustrate write data flow in a target subsystem. The Fig. 50 illustrates an R2T command flow, which is used by the target to inform the initiator that it is ready to accept a data write from the initiator. The initiator then sends the write which is received at the target and the internal data flow is illustrated in Fig. 51. The two figures together illustrate one R2T and data write pairs. Following are the major steps that are followed as
- 15 illustrated in Figs. 50 and 51 together:
1. The target host system in response to receiving a write request like that illustrated in Fig. 33, prepares the appropriate buffers to accept the write data and informs the storage flow controller when it is ready, to send the ready to transfer request to the initiator;
 - 20 2. The flow controller acknowledges the receipt of the request and the buffer pointers for DMA to the host driver;
 3. The flow controller then schedules the R2T command to be executed to the scheduler;
 4. The scheduler issues the command to one of the packet processor
 - 25 complexes that is ready to execute this command;
 5. The packet processor requests the session entry from the session cache controller;

6. The session entry is returned to the packet processor;
7. The packet processor forms a TCP packet and encapsulates the R2T command and sends it to the output queue;
8. The packet is then sent out to network media interface which then sends the
5 packet to the initiator. The security engine could be involved, if the transfer needed to be secure transfer;
9. Then as illustrated in Fig. 51, the initiator responds to R2T by sending the write data to the target. The network media interface receives the packet and queues it to the input queue;
10. The packet scheduler retrieves the packet from the input queue;
11. The packet is scheduled to the classification engine;
12. The classification engine provides the classified packet to the scheduler with the classification tag. The flow illustrated is for unencrypted packet and hence the security engine is not exercised;
13. The scheduler assigns the packet based on the flow based resource
15 assignment queue to packet processor queue. The packet is then transferred to the packet processor complex when the packet processor is ready to execute this packet;
14. The packet processor requests the session cache entry (if it does not already have it in its local cache);
15. The session entry is returned to the requesting packet processor;
16. The packet processor performs all the TCP/IP functions, updates the session entry and the storage engine extracts the PDU as the write command in response to the previous R2T. It updates the storage session entry and routes the packet to the host output queue for it to be transferred to the host buffer. The packet may be tagged with the memory
25 descriptor or the memory descriptor list that may be used to perform the DMA of this packet into the host allocated destination buffer; and

17. The host interface block performs the DMA, to complete this segment of the Write data command.

Fig. 52 illustrates the target read data flow. This flow is very similar to the initiator R2T and write data flow illustrated in Fig. 41. The major steps followed in this flow are as follows:

- 5 1. Input packet is received from the network media interface block;
2. Packet scheduler retrieves the packet from the input queue;
3. Packet is scheduled for classification;
4. Classified packet returns from the classifier with a classification tag;
 - a. Depending on the classification and flow based resource allocation,
- 10 the packet is assigned to a packet processor complex which operates on the packet
5. Packet processor complex looks-up session entry in the session cache (if not present locally);
6. Session cache entry is returned to the packet processor complex;
7. The packet processor determines the Read Command PDU and requests the
- 15 read data with a request to the flow controller;
8. The flow controller starts the DMA to the host interface;
9. Host interface performs the DMA and returns the data to the host input queue;
10. The packet processor complex receives the data from the host input queue;
- 20 11. The packet processor complex forms a valid PDU and packet around the data, updates the appropriate session entry and transfers the packet to the output queue; and

12. The packet is transferred to the output network media interface block which transmits the data packet to the destination.

The discussion above of the flows is an illustration of some the major flows involved in high bandwidth data transfers. There are several flows like fragmented data flow, error flows
5 with multiple different types of errors, name resolution service flow, address resolution flows, login and logout flows, and the like are not illustrated, but are supported by the IP processor of this invention.

The IP processor of this invention may be manufactured into hardware products in the chosen embodiment of various possible embodiments using a manufacturing process,
10 without limitation, broadly outlined below. The processor may be designed and verified at various levels of chip design abstractions like RTL level, circuit/schematic/gate level, layout level etc. for functionality, timing and other design and manufacturability constraints for specific target manufacturing process technology. The processor design at the appropriate physical/layout level may be used to create mask sets to be used for manufacturing the chip
15 in the target process technology. The mask sets are then used to build the processor chip through the steps used for the selected process technology. The processor chip then may go through testing/packaging process as appropriate to assure the quality of the manufactured processor product.

While the foregoing has been with reference to particular embodiments of the invention, it
20 will be appreciated by those skilled in the art that changes in these embodiments may be made without departing from the principles and spirit of the invention.

What is claimed is:

1. A hardware processor providing remote direct memory access capability on an IP network and using a TCP, SCTP or UDP protocol, or a combination of any of the foregoing, over IP networks.
- 5 2. A hardware processor providing remote direct memory access capability on an Ethernet network.
3. A hardware processor providing remote direct memory access capability on an IP network using a protocol selected from the group of protocols consisting of the group of protocols excluding TCP, SCTP and UDP.
- 10 4. A hardware processor bypassing the TCP/IP stack of a system and providing remote direct memory access capability.
5. The processor of claims 1 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS,
15 DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.
6. The processor of claim 5 having certain of its functions implemented in hardware and certain of its functions implemented in software.
7. The hardware processor of claim 6, said processor included as a companion
20 processor on a server chipset or a core logic chipset or an I/O chipset.
8. The combination of claim 7 wherein the server is a blade server, thin server, media server, streaming media server, appliance server, Unix server, Linux server, Windows or Windows derivative server, AIX server, clustered server, database server, grid computing server, VOIP server, wireless gateway server, security server, file server, network attached
25 storage server, or game server or a combination of any of the foregoing.
9. A hardware processor providing remote direct memory access capability for enabling data transfer and using a TCP, SCTP or UDP protocol or a combination of any of

the foregoing, said processor enabling storage and retrieval, to and from a storage system, of data transmitted over an IP network.

10. A hardware processor providing remote direct memory access capability for enabling data transfer and using a protocol selected from the group of protocols consisting of other than TCP, SCTP or UDP, said processor enabling storage and retrieval, to and from a storage system, of data transmitted over an IP network.

11. The hardware processor of claims 9 or 10 included as a companion processor on a server.

12. The processor of claim 11 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

13. The combination of claim 12 wherein said processor is embedded inside a chipset on the server's motherboard.

14. The hardware processor of claim 12 further comprising the hardware implemented function of data packet security.

15. The hardware processor of claim 12 further comprising the hardware implemented function of data packet scheduling.

16. The hardware processor of claim 12 further comprising the hardware implemented function of data packet classification.

17. The hardware processor of claim 12 included as a companion processor on a storage system's chipset.

18. The combination of claim 17 wherein said processor provides IP network storage capability for said storage system to operate in an IP based storage area network.

19. The combination of claim 18 further comprising a controller blade and said chipset is in an interface between said storage system and said storage area network.

20. The combination of claim 18 further comprising at least one additional storage system and means to access said at least one additional storage system and to control the storage function in said at least additional one storage system.

5 21. A hardware processor providing remote direct memory access capability for enabling data transfer and using a TCP or SCTP or UDP protocol or a combination of any of the foregoing over IP networks, said processor embedded in a server's host hardware components for providing IP networking capability.

10 22. A hardware processor providing remote direct memory access capability for enabling data transfer and using a protocol selected from the group of protocols consisting of other than TCP, SCTP or UDP over IP networks, said processor embedded in a server's host hardware components for providing IP networking capability.

15 23. The processor of claim 21 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

24. The processor of claim 23 wherein a subset of said hardware components is capable of accessing network storage to transmit or receive data to or from said storage over the Internet.

20 25. The processor of claim 23 wherein at least a subset of said hardware components is used as a blade in a blade server.

26. The processor of claim 23 wherein at least a subset of said hardware components is used as an adapter in a server.

25 27. The processor of claim 23 wherein said hardware components comprise a network connectivity for data transfer for a storage system.

28. A hardware processor providing remote direct memory access capability for enabling data transfer using a TCP, SCTP or UDP protocol or a combination of any of the foregoing, over IP networks, said processor included as part of a chipset of a host processor for providing offloading capability for said protocol.

29. A hardware processor providing remote direct memory access capability for enabling data transfer and using a protocol selected from the group of protocols consisting of other than TCP, SCTP or UDP over IP networks, said processor included as part of a chipset of a host processor for providing protocol offloading capability for said protocol.

5 30. The processor of claim 28 wherein said hardware processor is programmable and operates on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivatives, SGML, or HTML format of a combination of any of the foregoing.

10 31. The combination of claim 30 wherein said host processor forms part of an apparatus further comprising a high end server, workstation, personal computer, hand held device, router, switch or gateway capable of interfacing with wired or wireless networks, blade server, thin server, media server, streaming media server, appliance server, Unix server, Linux server, Windows or Windows derivative server, AIX server, clustered server,
15 database server, grid computing server, VOIP server, wireless gateway server, security server, file server, network attached storage server, or game server or a combination of any of the foregoing.

 32. The combination of claim 31 wherein at least one of said apparatus is a low power apparatus.

20 33. The combination of claim 31 wherein said processor is included within a microcontroller, a processor or a chipset of at least one of said apparatus.

 34. A hardware processor providing remote direct memory access capability for enabling data transfer using TCP over IP networks, said processor embedded in an IP storage area network switching system line card, said processor being programmable and
25 operating on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

 35. A switching system having a plurality of line cards, each said line card having
30 identification information based therein and comprising a hardware processor providing

remote direct memory access capability for enabling data transfer using TCP over IP networks, said processor being programmable and sending and receiving data packets also having identification information based therein, said packets transmitted, encapsulated or encoded using a iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, 5 FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format, or a combination of any of the foregoing.

36. The switching system of claim 35 wherein the processor on a first of said plurality of line cards compares said destination information of said data packets with said identification information of a plurality of said plurality of line cards and transmits said 10 packets to a second plurality of line cards based on said comparison.

37. The combination of claim 36 wherein said switching system provides multi-protocol support.

38. The combination of claim 36 wherein said switching system interfaces with an IP based storage area network and with a fibre channel, infiniband, serial ATA, SAS, IP 15 Storage, or Ethernet protocol, or a combination of any of the foregoing, using a secure or a non-secure mode of data transfer.

39. The combination of claim 38 wherein said IP Storage protocol is iSCSI, FCIP, iFCP, or mFCP or a combination of any of the foregoing.

40. The combination of claim 38 terminating traffic in a first of said protocols and 20 originating traffic in a second of said protocols.

41. A hardware processor providing remote direct memory access capability for enabling data transfer using TCP or SCTP or UDP over IP networks, said processor embedded in a chipset of a gateway controller of a storage area network.

42. A hardware processor providing remote direct memory access capability for 25 enabling data transfer and using a protocol other than a TCP or SCTP or UDP protocol over IP networks, said processor embedded in a chipset of a gateway controller of a storage area network.

43. The processor of claim 41 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP,

infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivatives, SGML, or HTML format, or a combination of any of the foregoing.

44. A hardware processor providing remote direct memory access capability for enabling data transfer of data traffic using TCP over IP networks, said processor embedded in a chipset of a storage system or a storage area network management appliance for enabling said appliance to transport TCP/IP packets in-band to said data traffic or out of band to said data traffic.

45. The hardware processor of claim 44 wherein said processor operates on said packets to apply an access control, intrusion detection, bandwidth monitoring, bandwidth management, traffic shaping, security, virus detection, anti-spam, quality of service, encryption, decryption, LUN masking, zoning, multi-pathing, link aggregation or virtualization function or policy or a combination of any of the foregoing.

46. A networking appliance comprising a hardware processor providing remote direct memory access capability for enabling data transfer from and to a data source, to and from a data destination, of data traffic transmitted, encapsulated or encoded using TCP over IP networks, said processor enabling said appliance to transport TCP/IP packets in-band to said data traffic or out of band to said data traffic.

47. The appliance of claim 46 wherein said processor operates on said packets to apply an access control, intrusion detection, bandwidth monitoring, bandwidth management, traffic shaping, security, virus detection, anti-spam, quality of service, encryption, decryption, LUN masking, zoning, multi-pathing, link aggregation or virtualization function or policy or a combination of any of the foregoing.

48. The appliance of claim 46 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

49. The combination of claim 46 wherein said hardware processor itself includes a processor for performing deep packet inspection and classification.

50. The combination of claim 49 wherein said hardware processor itself includes a processor for performing policy management or policy enforcement on a packet-by-packet basis.

5 51. The combination of claim 50 wherein said hardware processor performs a function of virtualization, policy based management, policy enforcement, operations in-band to said data traffic or operations out of band to said data traffic.

52. A hardware processor providing remote direct memory access capability for enabling data transfer using TCP, SCTP or UDP or a combination thereof over IP networks, said processor in at least one server in a cluster of servers.

10 53. A hardware processor providing remote direct memory access capability for enabling data transfer using a protocol other than TCP, SCTP and UDP over IP networks, said processor embedded in at least one server in a cluster of servers.

54. A hardware processor providing remote direct memory access capability on an Ethernet network, said processor embedded in at least one server in a cluster of servers.

15 55. The processor of claim 52 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

20 56. A chip set having embedded therein a hardware processor providing remote direct memory access capability for enabling data transfer using TCP, SCTP or UDP or a combination of any of the foregoing over IP networks.

25 57. The chip set of claim 56 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

58. The combination of claim 56 wherein said processor has certain of its functions implemented in hardware and certain of its functions implemented in software.

59. The combination of claim 57, said processor having a security engine or a classification engine, or a combination of said engines, said engines being on separate chips of said chip set.

60. A host processor having a mother board, said motherboard having thereon
5 one chip of a chip set, said one chip comprising a programmable hardware processor providing remote direct memory access capability for enabling data transfer using TCP, SCTP or UDP, or other session oriented protocol or a combination of any of the foregoing over IP networks.

61. A host having a mother board, said motherboard having thereon one chip of a
10 chip set, said one chip containing a programmable hardware processor providing remote direct memory access capability for enabling data transfer using TCP, SCTP or UDP or a combination of any of the foregoing over IP networks, said processor having input and output queues and a RDMA controller, said queues and controller being maintained on said host.

62. The combination of claim 60, further comprising input and output queues and
15 storage flow controller, said queues and said controller being implemented on a chip in said chipset other than a chip on a mother board of said host.

63. A multi-port hardware processor of a predetermined speed providing remote
20 direct memory access capability for enabling data transfer using TCP, SCTP or UDP or a combination of any of the foregoing over IP networks, said processor coupled to multiple input and output ports each having slower speed line rates than said predetermined speed, the sum of said slower line speeds being less than or equal to said predetermined speed.

64. The processor of claim 63 wherein said processor is programmable and
25 operates on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML protocol or a combination of any of the foregoing.

65. A hardware processor providing remote direct memory access capability for
30 enabling telecommunications or networking using TCP, SCTP or UDP or a combination of any of the foregoing over IP.

66. An integrated circuit hardware processor providing remote direct memory access (RDMA) capability, said processor for enabling data transfer using TCP over IP or SCTP over IP or UDP over IP or Ethernet networks, or a combination of any of the foregoing, to a host running an application, said hardware processor comprising:

- 5 a. registration circuitry for allowing said application to register a memory region of said host processor with said hardware processor for RDMA access;
- b. communication circuitry for exporting said registered memory region to at least one peer hardware processor having RDMA capability and for informing said peer of said host processor's desire to allow said peer to read data from or write data to said
10 registered memory region; and
- c. RDMA circuitry for allowing information transfer to and/or from said registered region of memory without substantial host processor intervention.

67. In a hardware processor providing remote direct memory access (RDMA) capability, said hardware processor circuit chip for enabling data transfer and using a TCP
15 over IP or SCTP over IP or UDP over IP or Ethernet networks, or a combination of any of the foregoing, the process of performing RDMA for an application running on a host processor, said process comprising:

- a. said application registering a region of memory of said host processor for RDMA;
- 20 b. said host processor making said region of memory available to a peer processor for access directly without substantial intervention by said host processor in said data transfer;
- c. said hardware processor communicating to said peer processor said host processor's desire to allow said peer processor to read data from or write data to said
25 region of memory; and
- d. said hardware processor enabling information transfer from or to said registered region of memory without substantial host processor intervention in said information transfer.

68. The process of claim 67 wherein said peer processor is part of a computing apparatus and said host processor and said computing apparatus are in a client to server, server to client, server to server, client to client or peer to peer session connection.

5 69. A hardware processor providing a transport layer remote direct memory access (RDMA) capability, said processor for enabling data transfer over a network using TCP over IP in one or more session connections, said processor including a TCP/IP stack, said stack including an interface to upper layer functions to transport data traffic, said stack providing at least one of the functions of:

- a. sending and receiving data, including upper layer data;
- 10 b. establishing transport sessions and session teardown functions;
- c. executing error handling functions;
- d. executing time-outs;
- e. executing retransmissions;
- f. executing segmenting and sequencing operations;
- 15 g. maintaining protocol information regarding said active transport sessions;
- h. maintaining TCP/IP state information for each of said one or more session connections;
- i. fragmenting and defragmenting data packets;
- 20 j. routing and forwarding data and control information;
- k. sending to and receiving from a peer, memory regions reserved for RDMA;
- l. recording said memory regions reserved for RDMA in an RDMA database and maintaining said database;
- 25 m. executing operations provided by RDMA capability;

70

- n. executing security management functions;
- o. executing policy management and enforcement functions;
- p. executing virtualization functions;
- q. communicating errors;
- 5 r. processing Layer 2 media access functions to receive and transmit data packets, validate the packets, handle errors, communicate errors and other Layer 2 functions;
- s. processing physical layer interface functions;
- t. executing TCP/IP checksum generation and verification functions;
- 10 u. processing Out of Order packet handling;
- v. CRC calculation functions;
- w. processing Direct Data Placement/Transfer;
- x. Upper Layer Framing functions;
- y. processing functions and interface to socket API's;
- 15 z. forming packet headers for TCP/IP for transmitted data and extraction of payload from received packets; and
- aa. processing header formation and payload extraction for Layer 2 protocols of data to be transmitted and received data packets; respectively.

20 70. The hardware processor of claim 69 wherein said data transfer includes storing and retrieving data.

71. The hardware processor of claims 69 or 70 having an IP layer interfacing with a media access layer to transport said data packets onto said network.

72. The hardware processor of claim 71 wherein said media access layer is wired or wireless Ethernet, MII, GMII, XGMII, XPF, XAUI, TBI, SONET, DSL, POS, POS-PHY, SPI

interface, SPI-4 or SPI derivative or other SPI derivative interface, Infiniband, or FC layer or a combination of any of the foregoing.

73. A hardware processor providing TCP or SCTP or UDP, or a combination of any of the foregoing, over IP operations including RDMA capability for data transfer over a network from or to an initiator and to or from a target, said operations requested by a host processor having a SCSI command layer and an iSCSI driver or an IP Storage driver, said hardware processor comprising:
- a. an RDMA mechanism;
 - b. a command scheduler for scheduling commands from the command layer of said host processor for operation in the processor;
 - c. first command queues for queuing commands from said host processor for existing sessions;
 - d. second command queues for queuing commands from said host processor for sessions that do not currently exist;
 - e. a database for recording the state of the session on which said command is transported, said database also for recording progress of RDMA for those of said commands that use RDMA;
 - f. a communication path between said processor and said SCSI layer for communicating status of command execution to said SCSI layer for processing; and
 - g. at least one transmit/receive engine and at least one command engine coupled to work together to interpret commands and perform appropriate operations for performing RDMA for storing/retrieving data to/from or transmitting/receiving data to/from said target or said initiator.

74. The combination of claim 73 wherein said at least one transmit/receive engine is implemented as a separate transmit engine and a separate receive engine.

75. The combination of claim 73 wherein said at least one transmit/receive engine and at least one command engine are implemented as at least one composite engine providing the transmit, receive and command engine functionality.

76. The combination of claim 73 wherein said first command queues are located partly in memory on said hardware processor and partly in memory off said hardware processor.

5 77. The combination of claim 73 wherein said second command queues are located partly in memory on said hardware processor and partly in memory off said hardware processor.

78. The combination of claim 76 wherein said memory off said hardware processor is memory included in a host system.

10 79. The combination of claim 77 wherein said memory off said hardware processor is memory included in a host system.

80. The combination of claim 76 wherein said memory off said hardware processor is RAM, DRAM, SDRAM, DDR SDRAM, RDRAM, FCRAM, FLASH, ROM, EPROM, EEPROM, QDR SRAM, QDR DRAM or other derivatives of static or dynamic random access memories or a combination of any of the foregoing.

15 81. The combination of claim 77 wherein said memory off said hardware processor is RAM, DRAM, SDRAM, DDR SDRAM, RDRAM, FCRAM, FLASH, ROM, EPROM, EEPROM, FLASH, ROM, EPROM, EEPROM, QDR SRAM, QDR DRAM or other derivatives of static or dynamic random access memories or a combination of any of the foregoing.

20 82. The combination of claim 76 wherein said memory on a chip not included in said host processor is located on a companion chip to said hardware processor.

83. The combination of claim 77 wherein said memory on a chip not included in said host processor is located on a companion chip to said hardware processor.

25 84. An iSCSI stack implemented on an integrated circuit chip providing remote direct memory access capability for use in transporting information over IP networks by transporting PDU's specified by the iSCSI standard in accordance with SCSI commands, said stack comprising hardware implementation of functions that are standard defined iSCSI functions, session establishment and teardown functions, or remote data memory access functions or a combination of any of the foregoing.

85. The iSCSI stack of claim 84 wherein said SCSI commands are queued in a command queue for execution by an IP Storage processor or an iSCSI processor.

86. An IP storage or iSCSI stack providing remote direct memory access capability for use in transporting information in active sessions or connections over IP
5 networks by transporting PDU's specified by the iSCSI standard or IP storage standard in accordance with SCSI commands, said stack providing an interface to the upper layer protocol functions in a host processor to transport storage data traffic, and performing at least one of the hardware implemented functions of

- a. sending and receiving data, including upper layer data;
- 10 b. sending and receiving command PDU's;
- c. establishing transport sessions and session teardown functions;
- d. executing error handling functions;
- e. executing time-outs;
- f. executing retransmissions;
- 15 g. executing segmenting and sequencing operations;
- h. maintaining protocol information regarding said active transport sessions;
- i. maintaining TCP/IP state information for each of said one or more session connections;
- 20 j. maintaining IP storage state information for each of said one or more session connections;
- k. fragmenting and defragmenting data packets;
- l. routing and forwarding data and control information;
- 25 m. sending to and receiving from a peer, memory regions reserved for RDMA;

- n. recording said memory regions reserved for RDMA in an RDMA database and maintaining said database;
- o. executing operations provided by RDMA capability;
- p. executing security management functions;
- 5 q. executing policy management and enforcement functions;
- r. executing virtualization functions;
- s. communicating errors;
- t. processing Layer 2 media access functions to receive and transmit data packets, validate the packets, handle errors, communicate errors and other Layer 2 functions;
- 10 u. processing physical layer interface functions;
- v. executing TCP/IP checksum generation and verification functions;
- w. processing Out of Order packet handling;
- x. CRC calculation functions;
- 15 y. processing Direct Data Placement;
- z. Upper Layer Framing functions;
- aa. processing functions and interface to networking socket API's;
- bb. processing functions and interface to IP storage or iSCSI driver and/or SCSI command layer;
- 20 cc. forming packet headers for TCP/IP for transmitted data and extraction of payload from received packets;
- dd. forming packet headers for IP storage or iSCSI PDUs for transmitted data and extraction of PDUs from received packets; and

ee. processing header formation and payload extraction for Layer 2 protocols of data to be transmitted and received data packets, respectively.

5 87. A TCP/IP stack providing transport layer remote direct memory access capability for use in transporting information in active sessions or connections over IP networks, said stack providing an interface to the upper layer protocol functions in a host processor to carry data traffic, and performing at least one of the hardware implemented functions of:

- a. sending and receiving data, including upper layer data;
- b. establishing transport sessions and session teardown functions;
- 10 c. executing error handling functions;
- d. executing time-outs;
- e. executing retransmissions;
- f. executing segmenting and sequencing operations;
- g. maintaining protocol information regarding said active transport
15 sessions;
- h. maintaining TCP/IP state information for each of said one or more session connections.
- i. fragmenting and defragmenting data packets;
- j. routing and forwarding data and control information;
- 20 k. sending to and receiving from a peer, memory regions reserved for RDMA;
- l. recording said memory regions reserved for RDMA in an RDMA database and maintaining said database;
- m. executing operations provided by RDMA capability;

- n. executing security management functions;
 - o. executing policy management and enforcement functions;
 - p. executing virtualization functions;
 - q. communicating errors;
 - 5 r. processing Layer 2 media access functions to receive and transmit data packets, validate the packets, handle errors, communicate errors and other Layer 2 functions;
 - s. processing physical layer interface functions;
 - t. executing TCP/IP checksum generation and verification functions;
 - 10 u. processing Out of Order packet handling;
 - v. CRC calculation functions;
 - w. processing Direct Data Placement/Transfer;
 - x. Upper Layer Framing functions;
 - y. processing functions and interface to socket API's;
 - 15 z. forming packet headers for TCP/IP for transmitted data and extraction of payload from received packets; and
 - aa. processing header formation and payload extraction for Layer 2 protocols of data to be transmitted and received data packets; respectively.
88. The TCP/IP stack of claim 69 further comprising memory for storing a
- 20 database to maintain various information regarding said active sessions or connections and TCP/IP state information for each of the sessions or connections.
89. The IP storage or iSCSI stack of claim 86 further comprising memory for storing a database to maintain various information regarding said active sessions or connections and state information for each of the sessions or connections.

90. The IP storage or iSCSI stack of claim 87 further comprising memory for storing a database to maintain various information regarding said active sessions or connections and TCP/IP state information for each of the sessions or connections.

5 91. The TCP/IP stack of claim 88 further comprising an interface that includes circuitry for interfacing to at least one layer that is a wired or wireless Ethernet, MII, GMII, XGMII, XPF, XAUI, TBI, SONET, DSL, POS, POS-PHY, SPI Interface, SPI-4 or other SPI derivative Interface, Infiniband, or FC layer.

10 92. The IP storage or iSCSI stack of claim 89 further comprising an interface that includes circuitry for interfacing to at least one layer that is a wired or wireless Ethernet, MII, GMII, XGMII, XPF, XAUI, TBI, SONET, DSL, POS, POS-PHY, SPI Interface, SPI-4 or other SPI derivative Interface, Infiniband, or FC layer.

15 93. The TCP/IP stack of claim 90 further comprising an interface that includes circuitry for interfacing to at least one layer that is a wired or wireless Ethernet, MII, GMII, XGMII, XPF, XAUI, TBI, SONET, DSL, POS, POS-PHY, SPI Interface, SPI-4 or other SPI derivative Interface, Infiniband, or FC layer.

20 94. A hardware implemented iSCSI or IP storage controller providing remote direct memory access useable in TCP or SCTP or UDP, or a combination of any of the foregoing over IP, said controller for transporting received iSCSI commands and PDUs, said controller having access to a data base for keeping track of data processing operations, said database being in memory on said controller, or in memory partly on said controller and partly in a computing apparatus other than said controller, said controller coupled to a host having a SCSI command layer and an iSCSI or IP storage driver, said controller having a transmit and a receive path for data flow, comprising:

25 a. a command scheduler for scheduling processing of commands, said scheduler capable of being coupled to said SCSI command layer and to said iSCSI or IP storage driver;

b. a receive path for data flow of received data and a transmit path for data flow of transmitted data;

c. at least one transmit engine for transmitting iSCSI or IP storage PDUs;

d. at least one transmit command engine capable of interpreting said PDUs and capable of performing operations comprising retrieving information from said host servicing remote direct memory access or iSCSI or IP storage commands or a combination thereof and keeping command flow information in said database updated as said retrieving progresses;

e. at least one receive command engine for receiving said iSCSI or IP storage commands; and

f. at least one receive command engine for interpreting said received iSCSI or IP storage or RDMA-commands or a combination thereof and capable of performing operations comprising storing or retrieving information to or from said host, servicing the received command and keeping command flow information in said database updated as storing or retrieving progresses.

95. An IP processor having RDMA capability, comprising an IP network application processor core for enabling TCP or SCTP or UDP or a combination of any of the foregoing over IP networks, comprising an intelligent flow controller, at least one packet processor, a programmable classification engine, a storage policy engine or network policy engine, a security processor, a session memory, a memory controller, a media interface and a host interface.

96. The combination of claim 94 wherein said at least one transmit and at least one receive engine is implemented as at least one composite engine providing transmit and receive functionality.

97. The combination of claim 94 wherein said at least one transmit and at least one receive engine and at least one command engine are implemented as at least one composite engine providing the transmit, receive and command engine functionality.

98. The IP processor of claim 95 wherein said host interface is a CSIX or an XPF, XAUI or an GMII, MII, XGMII, SPI, SPI-4 or other SPI derivative, 3GIO, PCI, PCI-Express, Infiniband, Fibre channel, RapidIO or Hypertransport type or a combination of any of the foregoing.

99. The IP processor of claim 95 further comprising a coprocessor for interfacing with a processor external to said IP processor.

100. The IP processor of claim 95 further comprising a system controller interface for connecting to a system controller.

5 101. The IP processor of claim 95 further comprising a control plane processor, said control plane processor functioning as a system controller.

102. An IP processor having RDMA capability for enabling TCP or SCTP or other session oriented protocols or UDP over IP networks, said processor comprising:

- a. an RDMA mechanism for performing RDMA data transfer;
 - 10 b. at least one packet processor for processing IP packets;
 - c. a session memory for storing IP session information;
 - d. at least one memory controller for controlling memory accesses;
 - e. a media interface for coupling to at least one network; and
 - f. a host interface for coupling to at least one host or a fabric interface
- 15 for coupling to a fabric.

103. The IP processor of claim 102 further comprising at least one of:

- a. an IP Storage session memory for storing IP session information;
- b. a classification processor for classifying IP packets;
- c. a flow controller for controlling data flow;
- 20 d. a policy processor for applying policies;
- e. a security processor for performing security operations;
- f. a controller for control plane processing;
- g. a packet scheduler;

- h. a packet memory for storing packets; or
- i. a combination of any of the foregoing.

104. The IP processor of claim 102 wherein any combination of said recited elements a-f or parts thereof are implemented in a single element.

- 5 105. The IP processor of claim 103 wherein any combination of said recited elements a-e or parts thereof are implemented in a single element.

106. A multiprocessor system comprising at least one data processor coupled to a plurality of IP processors for interfacing said at least one data processor to said IP processors, each having RDMA capability for enabling TCP or SCTP or other session oriented protocols or UDP over IP networks, said IP processor comprising:

- a. a RDMA mechanism for performing RDMA data transfer;
- b. at least one packet processor for processing IP packets;
- c. a session memory for storing IP session information;
- d. at least one memory controller for controlling memory accesses;
- 15 e. at least one media interface for coupling to at least one network; and
- f. a host interface for coupling to at least one host or fabric interface for coupling to a fabric.

107. The multiprocessor system of claim 106, said IP network application processor further comprising at least one of:

- 20 a. an IP Storage session memory for storing IP Storage session information;
- b. a classification processor for classifying IP packets;
- c. a flow controller for controlling data flow;
- d. a policy processor for applying policies;

- e. a security processor for performing security operations;
- f. a packet memory for storing packets;
- g. a controller for control plane processing;
- h. a packet scheduler;
- 5 i. a coprocessor interface for coupling to a peer processor; or
- j. a combination of any of the foregoing.

108. The multiprocessor system of claim 106 wherein two or more of said plurality of IP processors are coupled to each other.

109. The multiprocessor system of claim 107 wherein two or more of said plurality
10 of IP processors are coupled to each other.

110. The multiprocessor system of claim 108 wherein said two or more of said plurality of IP processors are coupled through a co-processor interface, or a host interface, or a bridge, or a combination of any of the foregoing.

111. The multiprocessor system of claim 109 wherein said two or more of said
15 plurality of IP processors are coupled through a co-processor interface, or a host interface, or a bridge, or a combination of any of the foregoing.

112. A TCP/IP processor engine having RDMA capability, said processor for processing Internet Protocol packets and comprising at least one of each of:

- a. an RDMA mechanism for performing RDMA data transfer;
- 20 b. a checksum hardware for performing checksum operations;
- c. a data memory for storing data used in the TCP/IP processor;
- d. an instruction memory for storing instructions used in the TCP/IP processor;
- e. an instruction fetch mechanism for fetching said instructions;

- f. an instruction decoder for decoding said instructions;
 - g. an instruction sequencer for sequencing said instructions;
 - h. a session database memory for storing TCP/IP session data; or
 - i. a session database memory controller for controlling said session
- 5 database memory;

or a combination of any of the foregoing items a through i; and

a host interface, or a fabric interface, or bus controller, or memory controller or combination of any of the foregoing for coupling to host or a fabric.

113. The TCP/IP processor engine of claim 112 further comprising at least one of:

- 10 a. a hash engine for performing hash functions;
- b. a sequencer manager for sequencing operations;
- c. a window operations manager for performing windowing operations to position packets within, and/or verify packets to be within, agreed windows;
- d. a classification tag interpreter for interpreting classification tags;
- 15 e. a frame controller for controlling data framing;
- f. an out of order manager for handling out of order data;
- g. a register file for storing data;
- h. a TCP state manager for managing TCP session states;
- i. a CRC component for performing CRC operations;
- 20 j. an execution resource unit or ALU for data processing;
- k. a TCP session database lookup engine for accessing session entries;
- l. an SNACK engine for selective negative acknowledgment;

- m. an SACK engine for selective positive acknowledgment;
- n. a segmentation controller for controlling the segmentation of data;
- o. a timer for event timing;
- p. a packet memory for storing packets; and
- 5 q. a combination of any of the foregoing.

114. An IP storage processor engine having RDMA capability, said processor for processing Internet Protocol packets and comprising at least one of each of:

- a. an RDMA mechanism for performing RDMA data storage;
- b. CRC hardware for performing CRC functions;
- 10 c. a data memory for storing data used in the processor;
- d. an instruction memory for storing instructions used in the processor;
- e. an instruction fetch mechanism for fetching said instructions;
- f. an instruction decoder for decoding said instructions;
- g. an instruction sequencer for sequencing said instructions;
- 15 h. an IP storage session database memory for storing IP storage session information;
- i. an IP storage session database memory controller for controlling said IP storage session database memory;
- j. a combination of any of the foregoing items a through j; and
- 20 k. a host interface, or a fabric interface, or bus controller, or memory controller or combination thereof for a host or to a fabric.

115. The IP storage processor engine of claim 114 further comprising at least one of:

- a. hash engine for performing hash operations;
- b. a sequencer manager for sequencing operations;
- c. a window operations manager for positioning packets within, and/or verifying received packets to be within, agreed windows;
- 5 d. a classification tag interpreter for interpreting classifications tags;
- e. an out of order manager for handling out of order data;
- f. a register file for storing data;
- g. a PDU storage classifier for classifying packets into various attributes;
- h. an IP storage state manager for managing IP storage session states;
- 10 i. a checksum component for performing checksum operations;
- j. an execution resource unit or ALU for data processing;
- k. a TCP session database lookup engine for accessing session entries;
- l. a SNACK engine for selective negative acknowledgment;
- m. a SACK engine for selective positive acknowledgment;
- 15 n. a segmentation controller for controlling the segmentation of data;
- o. a timer for event timing;
- p. a packet memory for storing packets; and
- q. a combination of any of the foregoing.

116. The combination of claim 113 wherein said register file stores packet
20 headers, pointers, contexts or session states or any combination thereof.

117. The combination of claim 115 wherein said register file stores packet
headers, pointers, contexts or session states or any combination thereof.

118. A processor for processing Internet data packets in one or more sessions and capable of executing transport layer RDMA functions, said processor including a session memory for storing frequently or recently used session information for a plurality of sessions.

5 119. A TCP/IP processor implemented in hardware and capable of implementing transport level RDMA functions, said processor including a session memory for storing session information for a plurality of sessions.

120. A processor for processing Internet data packets in one or more sessions, said processor comprising an RDMA mechanism, and a session memory for storing session information for a plurality of said sessions.

10 121. An IP storage processor implemented in hardware and capable of performing RDMA functions, and including a session memory for storing session information for a plurality of sessions.

122. A hardware implemented IP network application processor implementing remote direct memory access (RDMA) capability for providing TCP/IP operations in sessions
15 on information packets from or to an initiator and providing information packets to or from a target, comprising the combination of:

- a. data processing resources comprising at least one programmable packet processor for processing said packets;
- b. an RDMA mechanism capable of providing remote direct memory
20 access function between said initiator and said target;
- c. a TCP/IP session cache and memory controller for keeping track of the progress of, and memory useful in, said operations on said packets;
- d. a host interface controller capable of controlling an interface to a host computer in an initiator or target computer system or a fabric interface controller capable of
25 controlling an interface to a fabric; and
- e. a media independent interface capable of controlling an interface to the network media in an initiator or target.

123. The combination of claim 122 further comprising:

- a. at least one data processor for providing the functions of
 - i. packet classification for classifying said packets; or
 - ii. packet scheduling for scheduling operations on said packets to said data processing resources; or
- 5 b. at least one controller for controlling the flow of information in said IP network application processor; or
- c. a first queue capable of storing packets incoming from said initiator or target; and a second queue capable of storing packets outgoing to said initiator or target; or
- d. a host input queue and a host output queue capable of accepting data
10 to be sent from or providing data or packets received to the host processor or memory associated with the host processor of the initiator or the target; or
- e. a combination of any of the foregoing items a. through d.

124. The processor of claim 123 wherein said at least one data processor provides the function of security processing for performing security processes for said packets
15 classified as secure by said classification engine.

125. The processors of claim 122 wherein said initiator and said target are computing apparati capable of being in a session connection that is client to server, server to client, server to server, client to client or peer to peer or a combination of any of the foregoing.

20 126. The processors of claim 123 wherein said initiator and said target are computing apparati capable of being in a session connection that is client to server, server to client, server to server, client to client or peer to peer or a combination of any of the foregoing.

25 127. The processors of claim 124 wherein said initiator and said target are computing apparati capable of being in a session connection that is client to server, server to client, server to server, client to client or peer to peer or a combination of any of the foregoing.

128. The processor of claim 124 where said functions of packet classification, security processing and packet scheduling are implemented in independent data processors in said at least one data processor.

129. The processor of claim 122 wherein said at least one controller includes a host interface controller capable of providing access to said host.

130. The processor of claim 123 wherein said at least one controller includes a RDMA controller for controlling said operations, including the controlling of said RDMA mechanism.

131. The processor of claim 123 wherein said at least one controller includes a control plane processor.

132. The processor of claim 131 wherein said control plane processor is a system controller.

133. The processor of claim 131 wherein said at least one controller functions as a session controller to create new sessions, teardown existing sessions, and create session data base entries, and to remove, update, access or store said entries, in the session memory and the cache.

134. The process, in a hardware implemented control plane processor or session controller capable of executing transport level RDMA functions and coupled to a host processor or a remote peer, of creating new sessions and their corresponding session database entries responsive to new session connection requests received either from the host processor or the remote peer.

135. The process, in hardware implemented control plane processor or session controller capable of executing transport level RDMA functions and coupled to a host processor or a remote peer and including a TCP/IP hardware processor or an IP storage hardware packet processor, or a combination of any of the foregoing, of tearing down or removing sessions and their corresponding session database entries responsive to session connection closure requests received either from the host processor or the remote peer or as a result of the operation by the said TCP/IP packet processor or IP Storage packet processor or a combination of any of the foregoing.

136. The processor of claim 122 wherein said host interface controller interfacing with the said host interface is CSIX, XPF, XAUI, GMII, MII, XGMII, SPI, SPI-4 or other SPI derivative, PCI, PCI-Express, Infiniband, Fibre channel, RapidIO or Hypertransport interface or other bus for switching system fabric interfaces, or a combination of any of the foregoing.

5 137. The processor of claim 122 wherein said host interface type is a PCI, PCI-X, 3GIO, PCI-Express, Rapid IO or HyperTransport or other bus for a non-switching interface type, or a combination of any of the foregoing.

10 138. The processor of claim 122 wherein said media independent interface type is Ethernet, MII, GMII, XGMII, XAUI, XPF, TBI, Ethernet MAC/PHY, SONET, DSL, POS, POS-PHY, SPI, SPI-4 or other SPI derivative, Infiniband, or FC interface or a combination of any of the foregoing.

15 139. The process in a hardware implemented processor capable of remote direct memory access for enabling storage and retrieval of data in a host memory subsystem, in an initiator system or in a target system, where said data is transferred using one or more data packets over an IP network

- a. from or to said target system to and from a host in said initiator system, or
- b. from or to said initiator system to and from a host in said target system,
- 20 i. said initiator system and said target system each including at least one hardware implemented processor capable of enabling storage and retrieval of data packets over an IP network and said target system and said initiator system each having a connection to said IP network,
- (a) said process comprising providing a remote direct
- 25 memory access process for said storage and retrieval without said host substantially controlling said remote direct memory access capability.

140. The process of claim 139 wherein said remote direct memory access process includes:

a. registering a region of memory of said host of said target system for access by said initiator system;

b. informing said initiator of the identity of said region of memory;

c. making said region of memory directly available to said initiator
5 system for access; and

d. using said remote direct memory access capability, transferring said data to or from said region of memory without said host of said target system substantially controlling said remote direct memory access capability.

141. The process of claim 140 wherein said registering includes registering the
10 identity of said memory regions with said initiator system.

142. The process of claim 140 wherein said memory regions comprise RDMA buffers and said transferring includes enabling said RDMA buffers and allowing said initiator system to write/push or read/pull data to or from said RDMA buffers.

143. The process of claim 140 further comprising (a) said initiator supplying a
15 request for access to said memory regions, (b) said host performing address translation of said requested memory region and comparing it to said registered region, (c) said host performing security key verification and, if said comparing and said verification are successful, allowing data storage to or retrieval from said requested region without substantial host intervention.

20 144. The process of claim 140 further comprising (a) said initiator supplying a request for access to said memory region, (b) said hardware processor to said host performing address translation of said requested memory region and comparing it to said registered region of memory, (c) said companion processor performing security key verification and, if said comparing and said verification are successful, allowing data storage
25 to or from said requested memory region.

145. The process of claim 139 wherein said remote direct memory access process includes:

a. registering a region of memory of said host of said initiator system for access by said target system;

b. informing said target of the identity of said memory region;

5 c. making said memory region directly available to said target system for access; and

d. using said remote direct memory access capability, transferring said data to or from said memory region without said host of said initiator system substantially controlling said remote direct memory access capability.

10 146. The process of claim 145 wherein said registering includes registering the identity of said memory regions with said target system.

147. The process of claim 145 wherein said memory regions comprise RDMA buffers and said transferring step includes enabling said RDMA buffers and allowing said target system to push or pull data to or from said RDMA buffers.

15 148. The process of claim 145 further comprising (a) said target supplying a request for access to said memory regions, (b) said host performing address translation of said requested memory region and comparing it to said registered region, (c) said host performing security key verification and, if said comparing and said verification are successful, allowing data storage to or retrieval from said requested memory region without substantial host intervention.

20 149. The process of claim 145 further comprising (a) said target supplying a request for access to said memory region, (b) said hardware processor to said host performing address translation of said requested memory region and comparing it to said registered region of memory, (c) said companion processor performing security key verification and, if said comparing and said verification are successful, allowing data storage
25 to or from said requested memory region.

150. A peer system connected to a host system, each of said peer system and said host system comprising at least one hardware processor capable of executing a transport layer RDMA protocol on an IP Network.

151. The combination of claim 150 performing RDMA transfer from the peer system to the host system.

152. The combination of claim 150 performing RDMA transfer from the host system to the peer system.

5 153. The combination of claim 118 further comprising an interface capable of being connected to a media access control layer of a host processor.

154. The combination of claim 120 further comprising an interface capable of being connected to a media access control layer of a host processor.

10 155. The process of claim 144, said hardware implemented processor further comprising an interface to a media control layer of a host processor, said process further comprising processing incoming packets from said media access layer through full TCP/IP termination and deep packet inspection.

15 156. A hardware processor providing TCP or SCTP or other session oriented protocols, or UDP over IP or any combination of any of the foregoing, including RDMA capability for data transfer over a network from or to an initiator and to or from said target, said operations requested by a host processor, comprising:

- a. an RDMA mechanism;
- b. a command scheduler for scheduling commands or other operations from the command layer or socket layer or RDMA layer, or any combination of any of the foregoing, of said host processor for operation in the hardware processor;
- 20 c. first command queues for queuing commands or other operations from said host processor for existing sessions;
- d. second command queues for queuing commands or other operations from said host processor for sessions that do not currently exist;
- 25 e. a database for recording in database entries the state of the session on which said command or other operation or its associated data is transported, said database also for recording progress of RDMA for those of said commands or other operations that use RDMA; and

f. at least one transmit/receive engine and at least one command engine coupled together, said engines working together to interpret commands and perform appropriate operations for performing RDMA for storing/retrieving data to/from or transmitting/receiving data to/from said target or said initiator.

5 157. A hardware processor for enabling data transfer over IP networks, said processor embedded in a blade server for providing networking capability.

158. The hardware processor of claim 157 wherein said data transfer uses a TCP, SCTP or UDP, or other session oriented protocol or a combination of any of the foregoing.

10 159. The hardware processor of claim 122 wherein said data transfer uses a protocol selected from the group of protocols excluding TCP, SCTP and UDP.

160. The appliance of claim 46 wherein said processor operates on said packets to apply one or more policies on packets and said appliance is located in-band to said traffic to apply one or more policies on packets at substantially the full line rate.

15 161. The appliance of claim 46 wherein said processor operates on said packets to apply one or more policies on packets and said appliance is located out-of-band to said traffic.

20 162. The appliance of claim 160 wherein said appliance is coupled to the source and destination of said data transfer, said appliance collecting control and management information from said source and destination for use in controlling and managing said packets.

163. The appliance of claim 161 wherein said appliance is coupled to the source and destination of said data transfer, said appliance collecting control and management information from said source and destination for use in controlling and managing said packets.

25 164. The appliance of claim 160 wherein said one or more policies is access control, intrusion detection, bandwidth monitoring, bandwidth management, traffic shaping, security, virus detection, anti-spam, quality of service, encryption, decryption, LUN masking, multi-pathing, link aggregation, zoning, and virtualization

165. The appliance of claim 164 wherein said one or more policies is applied using deep packet inspection or packet header classification.

166. A cluster of servers, a plurality of said servers having a hardware processor having RDMA capability for enabling data transfer over IP networks.

5 167. A cluster of servers, a plurality of said servers having a hardware processor having RDMA capability for enabling data transfer over Ethernet.

168. A central processing unit running a plurality of applications, said central processing unit having a separate hardware processor having RDMA capability for enabling data transfer over IP networks, said processor for communication among said applications
10 running on said central processing unit.

169. The combination of claim 166 wherein said data transfer is accomplished using a TCP, SCTP or UDP, or other session oriented protocol or a combination of any of the foregoing.

170. The combination of claim 167 wherein said data transfer is accomplished
15 using a TCP, SCTP or UDP, or other session oriented protocol or a combination of any of the foregoing.

171. The combination of claims 166 wherein said data transfer is accomplished using a protocol selected from the group of protocols consisting of the group or protocols excluding TCP, SCTP and UDP.

20 172. The combination of claims 167 wherein said data transfer is accomplished using a protocol selected from the group of protocols consisting of the group or protocols excluding TCP, SCTP and UDP.

173. The host processor of claim 60 wherein said queues and said storage controller are maintained in software on said host processor.

25 174. The host processor of claim 60 including input and output queues and a storage flow controller, maintained partly in software on said host processor and partly in hardware on said programmable hardware processor.

175. The host processor of claim 60 including input and output queues and a storage flow controller, maintained partly in hardware on said host processor and partly in hardware on said programmable hardware processor.

5 176. The host processor of claim 60 including input and output queues and a storage flow controller, maintained partly in hardware and partly in software in said host processor, and partly in hardware on said programmable hardware processor.

177. The hardware implemented iSCSI controller of claim 94 wherein said receive path for data flow of received data and said transmit path for data flow of transmitted data comprise the same physical path used at different times.

10 178. A hardware processor providing transport layer RDMA capability for enabling telecommunications or networking over an IP network.

179. The hardware processor of claim 73, said hardware processor receiving iSCSI commands for transporting data from an initiator to a target or from a target to an initiator in an iSCSI session and storing the state of each said session in said database as
15 session entries.

180. The hardware processor of claim 73 wherein said commands are both dependent commands and unrelated commands, said dependent commands scheduled on said queues of the same connection and unrelated commands are scheduled on queues of different connections.

20 181. The hardware processor of claim 179 wherein said commands are both dependent commands and unrelated commands, said dependent commands scheduled on said queues of the same connection and unrelated commands are scheduled on queues of different connections.

182. The hardware processor of claim 73 where all commands are queued to the
25 same session connection.

183. The hardware processor of claim 73 supporting multiple connections per session connection.

184. The hardware processor of claim 73 wherein said database entries are carried as separate tables or carried together as a composite table and said entries are used to direct block data transfer over TCP/IP.

5 185. The hardware processor of claim 156 wherein said database entries are stored as fields and said fields are updated as the connection progresses through multiple states during the course of data transfer.

186. The hardware processor of claim 73 wherein said first command queues are located partly in memory on said hardware processor and partly in memory off said hardware processor.

10 187. The combination of claim 73 wherein said second command queues are located partly in memory on said hardware processor and partly in memory off said hardware processor.

15 188. For use in a hardware implemented IP network application processor having remote direct memory access capability and including an input queue and queue controller for accepting incoming data packets including new commands from multiple input ports and queuing them on an input packet queue for scheduling and further processing, the process comprising: accepting incoming data packets from one or more input ports and queuing them on an input packet queue; and de-queuing said packets for scheduling and further packet processing.

20 189. The process of claim 188 comprising marking ones of said packets to allow different policies to be applied to different ones of said packets, said different policies based on characteristics selected from the group of characteristics consisting of (1) speed of the one of said one or more ports from which a data packet is accepted, (2) network interface of said one port, and (3) priority assigned to said one port using a fixed or programmable
25 priority.

190. The process of claim 188 comprising transmitting said dequeued packets to data processing apparatus for classification by a classifier and for scheduling by a scheduler.

191. The process of claim 190 wherein said classification allows deployment of policies enforceable per packet, or per transaction, or per flow or per command boundary or per a command set or a combination of any of the foregoing.

5 192. The process of claim 190 wherein said scheduling schedules new ones of said classified commands and data packets for operation in said hardware implemented IP network application processor.

193. The process of claim 192 wherein said operation is performed by a packet processing resource.

10 194. The process of claim 193 comprising decrypting said packets that are classified as encrypted, and scheduling the decrypted packets for further classification and scheduling by a scheduler.

15 195. The process of claim 188 comprising detecting that some of said incoming packets are fragmented and storing fragments of said fragmented incoming packets in memory selected from the group of memory consisting of memory on said network application processor or memory external to said network application processor.

196. The process of claim 188 further comprising detecting that downstream resources are backed-up and storing the incoming data packets in memory on said network application processor or memory external to said network application processor.

20 197. The process of claim 195 further comprising detecting the arrival of all fragments of a fragmented one of said incoming packets and merging said fragments to form a complete incoming packet said stored fragments retrieved from memory on said network application processor or memory external to said network application processor.

25 198. The process of claim 196 further comprising detecting that downstream resources are able to accept packets and, responsive thereto, retrieving said stored packets from memory on said network application processor or memory external to said network application processor.

199. The process of claim 188 further comprising assigning a packet tag to each said incoming data packet and using said tag to control the flow of said incoming packets through said network application processor.

200. A packet scheduler and sequencer for scheduling to a classification engine and to additional data processor resources in a hardware IP processor (1) data packets Incoming to said IP processor over an IP network and (2) tasks relating thereto, comprising:

- a. a classification controller and scheduler for retrieving packet headers
5 from a queue controller and transmitting data said packets to said classification engine for determining classification and managing execution of said packets;
- b. a queue for receiving headers of said packets for transmission to said classification controller and scheduler;
- c. an input queue for transmitting commands, including security
10 commands, to said data processor resources for execution for fragmented packets or secure packets;
- d. a state controller and sequencer for receiving state information of processes active inside said hardware processor and including means for providing instructions for the next set of processes to be active;
- e. an interface to a memory controller for queuing said packets that are
15 fragmented packets or that are not processed due to resources being backed up;
- f. a resource allocation table for assigning received packets or commands to ones of said processor resources based on the current state of said resources;
- g. a packet memory store located within or external to the packet
20 scheduler and sequencer for storing packets, packet tags, or classification results;
- h. a priority selector and a packet fetch and command controller,
 - i. said priority selector for retrieving commands and packet tags
from respective queues based on assigned priority; and
 - 25 ii. said packet fetch and command controller for retrieving the packet tags and classification results from said packet memory store and scheduling the packet transfer to appropriate resources; and

i. storage for receiving such classification results and transmitting said packets to said data processor resources based on the said classification results.

201. A classification resource for classifying, and a packet scheduler and sequencer for scheduling, to data processor resources in a hardware IP processor (1) data packets incoming to said IP processor over an IP network and (2) tasks relating thereto, comprising:

a. a classification controller and scheduler for retrieving packet headers or packets and classification tags from a classified queue of the packets that are classified by the classification engine and scheduling for a state controller and sequencer to assign them for execution to the said data processor resources;

b. an input queue for transmitting commands, including security commands, to said data processor resources for execution for fragmented packets or secure packets;

c. a state controller and sequencer for receiving state information of processes active inside said hardware IP processor and including means for providing instructions for the next set of processes to be active inside said hardware IP processor;

d. an interface to a memory controller for queuing said packets that are fragmented packets or that are not processed due to resources being backed up;

e. a resource allocation table for assigning received packets to ones of said processor resources based on the current state of said resources;

f. a packet memory store located within or external to the packet scheduler and sequencer for storing packets, packet tags, or classification results;

g. a priority selector and a packet fetch and command controller,

i. said priority selector for retrieving commands and packet tags from respective queues based on assigned priority;

ii. said packet fetch and command controller for retrieving the packet tags and classification results from said packet memory store and scheduling the packet transfer to appropriate resources; and

h. storage for receiving such classification results and transmitting said packets to said data processor resources based on the said classification results.

202. The packet scheduler and sequencer of claim 200 further comprising a queue for receiving and storing substantially complete packets for forwarding to said data processing resources for deep packet inspection and processing.

203. The packet scheduler and sequencer of claim 201 further comprising a queue for receiving and storing substantially complete packets for forwarding to said data processing resources for deep packet inspection and processing.

204. The packet scheduler and sequencer of claim 200 wherein said classification controller and scheduler receives fragmented packets and secure packets, forwards said fragmented blocks to a defragmenter to assemble completed packets, and forwards said secure packets for security scheduling to a security engine.

205. The packet scheduler and sequencer of claim 201 wherein said classification controller and scheduler receives fragmented packets and secure packets, forwards said fragmented blocks to a defragmenter to assemble completed packets, and forwards said secure packets for security scheduling to a security engine.

206. The packet scheduler and sequencer of claim 200 wherein said security scheduling includes assigning at least one security algorithm to be performed on said secure packets by a security engine.

207. The packet scheduler and sequencer of claim 201 wherein said security scheduling includes assigning at least one security algorithm to be performed on said secure packets by a security engine.

208. The process of scheduling and sequencing Internet Protocol packets, including packet headers, and tasks to a classification engine and other execution resources of a hardware processor having RDMA capability and balancing workload to said resources comprising:

retrieving the packets and packet headers from a header queue and transmitting said headers to said classification engine;

receiving classification results from said classification engine and storing them to a classifier queue; and

managing the execution of said packets through said execution resources.

209. The process of scheduling and sequencing in a hardware processor
5 classifying Internet Protocol packets, said packets further comprising a classification tag and packet descriptor, to and through execution resources of a hardware processor having RDMA capability and balancing workload to said resources comprising:

retrieving the classified packets, classification tag and results, from a classified queue receiving these from the classification engine; and

10 managing the execution of said packets through the execution resources.

210. The process of claim 202 further comprising detecting that deep packet inspection is required, and in response to said detecting, transmitting the complete packets to said classification for routing to said scheduler after classification.

211. The process of claim 200 further comprising detecting fragmented packets
15 and, in response to said detecting, defragmenting said fragmented packets.

212. The process of claim 208 further comprising detecting fragmented packets and, in response to said detecting, defragmenting said fragmented packets.

213. The process of claim 209 further comprising detecting secure packets and scheduling said secure packets to a security engine.

20 214. The process of claim 210 further comprising detecting secure packets and scheduling said secure packets to a security engine.

215. The process of claim 209 further comprising a scheduler state controller and sequencer receiving state information of various packet transactions and operations active
25 inside said hardware processor and providing to said execution resources instructions for subsequent packet transactions and operations.

216. The process of claim 210 further comprising a scheduler state controller and sequencer receiving state information of various packet transactions and operations active

inside said hardware processor and providing to said execution resources instructions for subsequent packet transactions and operations.

217. The process of claim 203 further comprising said scheduler retrieving said packets from an input packet queue and scheduling these packets in a queue of a predetermined execution resource depending on said classification results received from said classifier.

218. The process of claim 208 wherein said managing includes directing the packets to be stored in a packet memory store.

219. The process of claim 209 wherein said managing includes directing the packets to be stored in a packet memory store.

220. The process of claim 214 further comprising retrieving said packets and/or said classification tag including the packet descriptor or identifier when either is scheduled for operation.

221. The process of claim 213 further comprising said state controller and sequencer identifying the execution resource that should receive the packet for operation, creating a command including a packet classification tag identifying said packet and assigning said created command to the appropriate one of said data processing resources.

222. The process of claim 219 further comprising a priority selector retrieving said command and packet tag from the respective queues based on the assigned priority of said packet in a resource entry in a resource allocation table and passing said commands and packet tags to a packet fetch and command controller.

223. The process of claim 221 further comprising retrieving the packet from the packet memory store along with said classification results and scheduling the packet for transfer to the appropriate resource when the said resource is ready to accept a new packet for processing.

224. The process of claim 222 further comprising retrieving the packet classification tag including the packet identifier or descriptor and scheduling said tag for transfer to the appropriate data processing resource when said resource is ready to accept a new packet for processing.

225. The process of claim 223 further comprising a bus interface of the respective receiving data processing resource interpreting said command and accepting the packet and classification tag for operation.

5 226. The process of claim 224 further comprising said data processing resource requesting the said packet scheduler and sequencer's packet fetch and command controller to the retrieval of the entire packet.

227. The process of claim 225 further comprising retrieving the packet from the packet memory store and transporting it to the data processing resource.

10 228. The process of claim 226 further comprising retrieving the packet from the packet memory store and transporting it to the data processing resource.

229. The process of claim 226 further comprising a bus interface of the respective receiving data processing resource interpreting said command and accepting the packet and the classification tag for operation.

15 230. The process of claim 226 further comprising an execution engine informing said scheduler when the packet operation is complete; and

said scheduler and a retirement engine retiring the packet from its state when the packet is scheduled for its end destination, thereby freeing up said resource entry in said resource allocation table.

20 231. The process of claim 217 further comprising an execution engine informing said scheduler when the packet operation is complete; and

said scheduler and a retirement engine retiring the packet from its state when the packet is scheduled for its end destination, thereby freeing up said resource entry for said packet in said resource allocation table.

25 232. The process of claim 221 further comprising using the resource allocation table to assign the received packets to specific resources, depending on the current internal state of said resources.

233. The process of claim 230 further comprising detecting packets that belong to the same flow or connection and/or are dependent on an ordered execution and assigning

said detected packets to the same data processing resource by using a current session database cache state entry buffered in said same data processing resource without retrieving new session database entries external to said data processing resource.

234. The process of claim 232 further comprising a memory controller queuing
5 packets that are fragmented packets for defragmenting; or queuing complete packets for the case in which the scheduler queue is backed-up due to a packet processing bottleneck downstream.

235. A hardware data processing classifier engine for classifying Internet data
10 packets for traverse through a utilizing system, said classifying being accomplished according to the type of the packet, the protocol type, the port addresses of the packet, the source of the packet or the destination of the packet or a combination of any of the foregoing.

236. The classifier engine of claim 235 wherein said packet is in an iSCSI, iFCP,
15 infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, HTML TCP, SCTP or UDP format or a combination of any of the foregoing;

said classifier engine capable of testing a single packet field or a plurality of packet fields; and

20 including at least one input queue or at least one packet interface and input packet buffer for receiving input Internet Protocol packets for classification from a scheduler.

237. The classifier engine of claim 235 wherein said packet is in an iSCSI, iFCP,
infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS,
DAFS, HTTP, XML, XML derivative, SGML, HTML TCP, SCTP or UDP format or a combination of any of the foregoing;

25 said classifier engine capable of testing a single packet field or a plurality of fields; and

including at least one packet interface and at least one input packet buffer for receiving input Internet protocol packets for classification from the input queue and controller

238. The classifier engine of claim 235 that is a content addressable memory based classifier engine or a programmable classifier engine.

239. The process of a hardware classifier engine, in a hardware IP processor having a memory for storing a database of IP session entries and execution resources,
5 classifying Internet Protocol packets in accordance with an attribute including examining fields of a received packet to identify the type of the packet, the protocol type, the packet's port addresses, the packet source and the packet destination.

240. The process of claim 239 comprising:

obtaining input packets, including packet headers, from a scheduler;
10 queuing said packets and/or packet headers for classification;
fetching the next available packet in the queue by a packet classification sequencer and extracting the appropriate packet fields based on predetermined global packet field descriptor sets;
15 passing said fields to a memory or ALU to perform said classification, matching of said fields to programmed values to identify the next set of fields to be compared;
in response to said matching collecting action tags or event tags in a result compiler; and
acting on said matching to update data in said memory associated
20 with a specific condition or rule match.

241. The process of claim 239 comprising:

obtaining input packets, including packet headers, from a input queue and controller;
25 queuing said packets and/or packet headers for classification in a second queue;

fetching the next available packet in the second queue by a packet classification sequencer and extracting the appropriate packet fields based on predetermined global packet field descriptor sets;

5 passing said fields to a memory or ALU to perform said classification, and matching of said fields to programmed values to identify the next set of fields to be compared;

 in response to said matching, collecting action tags or event tags in a result compiler; and

10 acting on said matching to update data in said memory associated with a specific condition or rule match.

242. The process of claim 240 wherein said memory is a content addressable memory array, said content addressable memory array being programmed with said packet fields, their expected values and the corresponding one of said database entries programmed with the action on said matching said action tags, including the next packet field to compare.

15

243. The process of claim 241 wherein said memory is a content addressable memory array, said content addressable memory array being programmed with said packet fields, their expected values and the corresponding one of said database entries programmed with the action on said matching said action tags, including the next packet field to compare.

20

244. The process of classifying of claim 242, further comprising programming said content addressable memory array through initializing said database, said database accessible for programming through a system interface.

245. The process of claim 244 further comprising providing said system interface with a connection to a host or a control plane processor.

25

246. The classification process of claim 239 further comprising generating, after completing classification, a classification tag that identifies the path to be traversed by said packet within said utilizing system.

247. The classification process of claim 240 further comprising generating, after completing classification, a classification tag that identifies the path to be traversed by said packet within said utilizing system.

5 248. The classification process of claim 241 further comprising generating, after completing classification, a classification tag that identifies the path to be traversed by said packet within said utilizing system.

249. The classification process of claim 247, wherein said classification tag is usable by a plurality of said execution resources to avoid performing the same classification tasks more than once.

10 250. The classification process of claim 248 wherein resulting classification tags are usable by a plurality of said execution resources to avoid performing the same classification tasks more than once.

15 251. The classification process of claim 239 further comprising providing a classifier retirement queue for holding the packets that are classified and are waiting to be retrieved by a scheduler.

252. The classification process of claim 246 further comprising providing a classifier retirement queue for holding the packets that are classified and are waiting to be retrieved by a scheduler.

20 253. The classification process of claim 249 further comprising providing a classification database for storing fields, their expected values and actions on match for said packets, said database being extendible using a database extension interface.

254. The classification process of claim 250 further comprising providing a classification database for storing fields, their expected values and actions on match for said packets, said database being extendible using a database extension interface.

25 255. The classification process of claim 251 wherein said extension interface is implemented using pipeline control logic.

256. The classification process of claim 246 further comprising providing an action interpreter, an ALU or execution resource and a range matching component to provide

capabilities to allow programming (a) storage policies in the form of rule action tables, (b) network policies in the form of rule action tables, and (c) required actions if certain of said network policies or storage policies or both are met.

257. The classification process of claim 256 comprising programming and
5 collecting said policies or rule action tables, or a combination of any of the foregoing, through a host interface.

258. An Internet Protocol packet processor including RDMA capability and capable of being coupled to a scheduler, to a TCP engine or a storage engine, or a combination of any of the foregoing, said processor comprising:

10 an instruction decoder and sequencer;

an instruction memory or data memory or a combination of any of the foregoing;

an execution resource;

a bus controller or memory controller or a combination of any of the foregoing;
15 said Processor comprising:

fetching instructions from said instruction memory;

decoding them and sequencing them through said execution resource by said instruction decoder and sequencer;

transmitting said packets from said scheduler to said bus controller;
20 using said bus controller for moving the data packets from said scheduler to said data memory for operation and/or for moving said data packets to and/or from said TCP engine and to and from said storage engine for processing said packet by:

(a) extracting data;
25 (b) generating new packets in response to said packet processing code; or

108

(c) transferring said extracted data or newly generated packets or combination thereof to an output interface;

said extracting, generating and transferring enabling transmission to a media interface or a host processor interface.

5

259. A programmable TCP/IP processor engine, said processor having RDMA capability and used for processing Internet Protocol packets, said TCP/IP processor engine comprising:

10

a checksum component for TCP/IP checksum verification and for new checksum generation;

a data memory for storing said packets;

an execution resource;

15

a packet look-up interface for assisting an execution resource and an instruction sequencer for providing access to said data packets or predetermined data packet fields thereof;

an instruction decoder to direct said TCP/IP processor engine operation based on the results of a classification processor;

20

a sequence and window operation manager providing specific segmenting, sequencing and windowing operations for use in TCP/IP data sequencing calculations;

and further comprising:

25

a hash engine used to perform hash operations against predetermined fields of the packet to perform a hash table walk to determine the correct session entry for said packet;

a register file for extracting predetermined header fields from said packets for TCP processing;

109

pointer registers for indicating data source and destination;

context register sets for holding multiple contexts for packet execution;

5 said multiple contexts allowing, in response to a given packet execution stalling, another context to be invoked to enable said TCP/IP processor engine to continue the execution of another packet stream;

10 said TCP/IP processor engine having a cache for holding recently or frequently used session entries, including connection IDs, for local use;

and further having an interface for informing a packet scheduler of the connection IDs that are cached for each TCP/IP processor engine resource.

15 260. The TCP/IP processor engine of claim 259 further comprising a session database lookup engine and a session manager which, in response to an indication that the packet processor does not hold the session entry for the specific connection required for said session, work with said hash engine to retrieve said session entry from a global session memory through a memory controller interface and to replace said session entry in said
20 packet processor.

25 261. The combination of claim 260 wherein the session manager includes means, operative upon fetching of a session entry or its fields corresponding to a packet being processed by said TCP/IP packet processor, said means working with said hash engine to generate a session identifier for retrieving the corresponding session entry or its fields from a session database cache.

262. The combination of claim 260 wherein the session manager includes means, operative upon storing of a session entry or its fields corresponding to the packet being processed by said TCP/IP processor, said means acting with the said hash engine to

generate a session identifier, for storing the corresponding session entry or its fields to the session database cache to replace a session entry as a result of packet processing.

263. The session database look-up engine of claim 260 further comprising means operative upon the fetching of a new session entry from said global session memory, said
5 means storing the replaced session entry in said global session memory.

264. The TCP/IP processor engine of claim 260 wherein said session entry in said cache is exclusively cached to one of a plurality of processors so that request for access to said cache by more than one of said plurality does not cause any race conditions by non-exclusive access.

10 265. The TCP/IP processor engine of claim 264 wherein when a session entry is exclusively cached to one processor, and another processor requests said session entry, said entry is transferred to the requesting processor with exclusive access to said session entry.

15 266. The TCP/IP processor engine of claim 259 further comprising a TCP state machine capable of state transitions and capable of having a current state and generating a next state, said state machine receiving

a state information stored in the session entry; and

appropriate fields affecting said state information from a fetched or newly received packet being processed for allowing the state machine to generate said next state if
20 there is a state transition, and updating information in said session entry in cache to indicate said next state.

267. The TCP/IP processor engine of claim 260 further comprising a programmable frame controller and out of order manager used for extracting frame information from said packets and performing operations for execution of packets received
25 out of order from the expected sequence for their session or flow .

268. The TCP/IP processor engine of claim 267 operating on an upper layer and having an upper layer framing mechanism used by said programmable frame controller and out of order manager to extract the embedded PDUs from packets arriving out of order and allowing said PDUs to be directed to an end buffer destination.

269. The TCP/IP processor engine of claim 268 wherein said programmable frame controller operates on retransmitted packets.

270. The combination of claim 268 wherein the frame controller and out of order manager includes a cyclic redundancy check generator for identifying, verifying and
5 delineating markers in the upper layer frames from the received network packets or generating the upper layer frame markers using CRC codes for packets directed to the network.

271. The TCP/IP processor engine of claim 259 including a storage engine queue which, if a packet belongs to a storage data transfer, receives one of said packets having a
10 storage payload, or the storage payload of said packet, for processing by a storage engine.

272. The TCP/IP processor engine of claim 259 including a segmentation controller for segmenting data to be sent on the IP Network to create valid packets to transport said segmented data on the IP network.

273. The TCP/IP processor engine of claim 259 including a DMA engine for
15 retrieving packets or commands or combinations thereof from a scheduler or a host, and storing said packets or commands or data to internal memory of the packet processor for further processing by the packet processor.

274. The TCP/IP processor engine of claim 273 including a processor for
20 processing the said packets having packet data and extracting the packet data for transfer to an end buffer destination in a host processor.

275. The TCP/IP processor engine of claim 273 including at least one processor for processing the said host commands or retrieving outgoing host data or a combination of any of the foregoing using the said DMA engine, for additional processing by the TCP/IP processor engine to form an outgoing packet for transfer onto the IP network.

276. The TCP/IP processor engine of claim 259 further comprising means in said
25 packet processor engine for additional processing of said packet by execution of additional processing application code.

277. The TCP/IP processor engine of claim 276 wherein said additional processing application code can be executed on said packet before or after processing by the TCP/IP processor engine or an IP Storage processor engine or a combination thereof.

278. The process, using a hardware TCP/IP processor engine capable of
5 executing a transport layer RDMA protocol, of transferring data from a host processor to an output media interface of said host system using said RDMA protocol, said process including said TCP/IP processor engine forming headers and forming segmentation for said data to be transferred.

279. The process of claim 278 wherein at least one of said headers is a header
10 used for transporting data through a storage area network, a local area network, a metro area network, a system area network, a wireless local area network, a personal area network, a home network, a wide area network or a combination of any of the foregoing.

280. A packet scheduler and sequencer for scheduling to data processor
resources in a hardware IP processor (1) data or commands incoming to said IP processor
15 from a host processor, and (2) tasks relating thereto, comprising:

- a. a host command queue for queuing incoming host commands;
- b. said host command queue for transmitting commands, including security commands, to said data processor resources for execution;
- c. a state controller and sequencer for receiving state information of
20 processes active inside said hardware processor and including means for providing instructions for the next set of processes to be active;
- d. a priority selector and a packet fetch and command controller,
 - i. said priority selector for retrieving commands and packet tags
from respective queues based on assigned priority, and
 - 25 ii. said packet fetch and command controller for retrieving the commands from the said command queue or packet tags from said packet memory store and classification results and scheduling the packet or command transfer to appropriate resources; and

e. a resource allocation table for assigning received packets or commands to ones of said processor resources based on the current state of said resources.

5 281. The packet scheduler and sequencer of claim 280, wherein said packet fetch and command controller is for retrieving the packet tags, data or the entire packet from said packet memory.

282. The TCP/IP processor engine of claim 260 wherein said session entry in cache is cached using the Modified Exclusive Shared Invalid algorithm.

10 283. A programmable IP Storage processor engine said processor having RDMA capability, and used for processing IP Storage Protocol Data Units transported over an IP network using Internet Protocol packets having fields including PDUs with fields, said IP Storage processor comprising:

a cyclic redundancy check component for IP Storage PDU CRC verification or for new CRC generation or a combination of any of the foregoing;

15 a data memory for storing said packets and/or said protocol data units;

a packet look-up interface for providing access to said packets or PDUs or packet fields or PDU fields or a combination of any of the foregoing;

an execution resource;

20 an instruction decoder and an IP Storage PDU classifier to direct IP Storage processor engine operation;

a sequence manager providing specific sequence operations for use in IP Storage data sequencing calculations;

and further comprising:

25 an IP Storage session manager including a hash engine to perform hash operations against predetermined fields of the packet to perform a hash table walk to determine the correct session entry for said packet;

114

a register file for extracting predetermined header fields from said packets and PDUs for IP Storage processing;

pointer registers for indicating data source and destination; and

context register sets for holding multiple contexts for packet execution;

said multiple contexts allowing, in response to a given packet execution stalling, another context to be invoked to enable said IP Storage processor engine to continue the execution of another packet stream;

said IP Storage processor engine having a cache for holding recently or frequently used session entries, including connection IDs, for local use and further having an interface for informing a packet scheduler of the connection IDs that are cached for each IP Storage processor engine resource.

284. The IP Storage processor engine of claim 283 further comprising an IP Storage session database lookup engine and an IP Storage session manager which, in response to an indication that the packet processor does not hold the IP Storage session entry for a specific connection for said session, acts with said hash engine to retrieve said IP Storage session entry from a global session memory through a memory controller interface and to replace an IP storage session entry in said packet processor.

285. The combination of claim 284 wherein the IP storage session manager further comprises means operative upon fetching of a session entry or its fields corresponding to the packet being processed by the IP Storage processor engine, said means working with the said the hash engine to generate the session identifier, for retrieving the corresponding session entry or its fields from said cache.

286. The combination of claim 284 wherein the session manager further comprises means operative upon storing of a session entry or its fields corresponding to the packet

being processed by the IP Storage processor engine, said means working with the said classifier and hash index to generate the session identifier, for storing the corresponding session entry or its fields to the said cache to reflect updates to those fields or entries as a result of packet processing.

5 287. The IP Storage session database look-up engine of claim 284 further comprising means operative upon the fetching of a new session entry from said global session memory, said means for storing the replaced session entry in said global session memory.

10 288. The IP Storage processor engine of claim 283 wherein a session entry in said cache is exclusively cached to one of a plurality of processors so that request for access to said cache by more than one of said plurality does not cause any race conditions by non-exclusive access.

15 289. The IP Storage processor engine of claim 288 wherein when a session entry is exclusively cached to one processor, and another processor requests said session entry, said entry is transferred to the requesting processor with exclusive access to said session entry.

20 290. The IP Storage processor engine of claim 288 further comprising a IP Storage state machine capable of state transitions and capable of having a current state and generating a next state, said state machine receiving
a state information stored in the session entry; and

appropriate fields affecting said state information from the fetched or newly received packet/PDU being processed for allowing the state machine to generate the next state if there is a state transition, and updating the next state information in said cache.

25 291. The IP Storage processor engine of claim 283 further comprising a segmentation controller for segmenting IP Storage data to be sent on the said IP Network to create valid packets for the TCP/IP packets that transport the IP Storage PDUs.

292. The IP storage processor engine of claim 291 having an upper layer framing mechanism used by a programmable frame controller/out of order manager to (a) extract the

PDU's from packets arriving out of order and allowing them to be directed to the end buffer destination, and (b) operating on retransmitted packets.

293. The combination of claim 292 wherein the frame controller/out of order manager further comprises a cyclic redundancy check generator for identifying, verifying and delineating markers in the upper layer frames from the received network packets or generating the upper layer frame markers using CRC codes for packets directed to the network.

294. The combination of claim 283 wherein said IP Storage processor engine comprises a TCP/IP Interface to couple to a TCP/IP engine, which interface (a) directs IP Storage PDU's transported using TCP/IP to the said TCP/IP engine and (b) receives IP Storage data PDU's extracted by the TCP/IP engine for processing by the IP Storage processor.

295. The combination of claim 283 wherein said IP Storage processor engine further comprises a DMA engine for retrieving packets from a scheduler or from a TCP/IP engine or commands received from a host, and storing said packets or host commands to internal memory of the packet processor for further processing by a packet processor.

296. The combination of claim 295 wherein said IP Storage processor engine processes the said packets and extracts therefrom the packet/PDU data or payload and transports it from the said packet to an end buffer destination in a host processor.

297. The combination of claim 295 wherein said IP Storage processor engine processes the said host commands or retrieves outgoing host data or a combination thereof using the said DMA engine for additional processing to form an outgoing packet/PDU for eventual transfer onto the IP Network.

298. The combination of claim 293 wherein said IP Storage processor engine further comprises processing resources in said processor engine for additional processing of said packet by execution of additional processing application code.

299. The combination of claim 298 wherein said IP Storage processor engine executes said additional processing application code on said packet before or after processing by the a TCP/IP processor engine and/or an IP storage processor engine.

300. The process, using an IP Storage processor capable of executing a transport layer RDMA protocol, of transporting data from a host processor to an output media interface of a host system, said process including said IP Storage processor forming headers and forming segmentation for said data to be transferred.

5 301. The process of claim 300 wherein at least one of said headers is a header used for transporting data through a storage area network, a local area network, a metro area network, a system area network, a system area network, a wireless local area network, a personal area network, a home network or a wide area network or a combination of any of the foregoing.

10 302. The combination of claim 283 wherein said IP Storage processor engine further comprises at least one Storage command execution engine to perform the operations for the IP Storage PDU command type decoded by the said instruction decoder and/or the said IP Storage PDU classifier.

15 303. The combination of claim 293 wherein said IP Storage processor engine further comprises an IP Storage initiator command decoder capable of interpreting commands received from the host processor to operate on the command working with the execution resources of the IP storage processor engine.

20 304. The combination of claim 283 wherein said IP Storage processor engine further comprises a IP Storage PDU creator that creates a PDU responsive to command decoding by the said command decoder.

305. The combination of claim 304 wherein said PDU creator creates the PDUs for protocol translation or virtualization operation directed by the said command decoder.

25 306. The combination of claim 283 wherein said IP storage processor engine further comprises means of transporting the said PDU to a TCP/IP processor engine for creation of appropriate network headers for transporting the packet to the network interface or the host interface as directed by the decoded command.

307. For use in a hardware implemented IP network application processor capable of executing transport level RDMA functions and having one or more output queues for accepting outgoing data packets, including new commands, from one or more packet

processor engines, TCP/IP processor engines or IP Storage processing engines, directing said packets on to an output port interface for sending them out to a network interface, said interface sending said packets on to the network, through one or more output packets, said process including: accepting incoming data packets from one or more packet processing
5 engines and queuing them on said output packet queue; and de-queuing said packets for delivery to said output port based on the destination port information in said packet.

308. The process of claim 307 further comprising decoding the packet tag of said packets to allow different policies to be applied to different ones of said packets, said different policies based on (1) speed of the one of said one or more ports to which a data
10 packet is scheduled, (2) network interface of said one port, and (3) priority assigned to said one port using a fixed or programmable priority.

309. The process of claim 307 further comprising decoding the packet tag of said packets to determine the need for application of security policy and directing the said packets needing the said security policy application to the clear packet port of a security
15 engine and later accepting the outgoing packets from the security engine with security policy applied.

310. The process of claim 307 comprising applying different policies to different ones of said packets, said different policies based on (1) need for encryption of the outgoing packet as directed by a packet tag or a packet handling command, (2) type of encryption
20 algorithm to be applied, (3) message authentication code or header creation, (4) type of authentication algorithm to be used or (5) a combination of one or more of the foregoing.

311. A storage flow and RDMA controller for controlling the flow of storage or non-storage data or commands, or a combination thereof, which may or may not require RDMA, for scheduling said commands and data to a scheduler or host processor or control plane
25 processor or additional data processor resources including one or more packet processors, in a hardware IP processor, wherein (1) data or commands are incoming to said IP processor over an IP network, or (2) data or commands are incoming from the host interface from the host processor, said storage flow and RDMA controller comprising at least one of:

a. a command scheduler, state controller and sequencer for retrieving
30 commands from a one or more command queues and sending the said command to the

control plane processor or the scheduler or one or more of said packet processors for further processing, and managing execution of the said commands by these resources;

b. a new commands input queue for receiving new commands from a host processor;

5 c. an active command queue for holding commands that are being processed, including newly scheduled commands processed from the said new commands queue;

d. an output completion queue for transmitting the status and the completed commands or their ID to the host processor for the host to take necessary
10 actions, including updating statistics related to the command and/or the connection, any error handling, releasing of any data buffers based on the command under execution;

e. an output new requests queue for transmitting to the host processor and the drivers running on the host, incoming commands from the packets received by the said IP packet processor for the host to take appropriate actions which may include
15 allocating appropriate buffers for receiving incoming data on the connection, or acting on RDMA commands, or error handling commands or any other incoming commands;

f. a command look-up engine to look-up the state of the commands being executed including look-up of associated memory descriptors or memory descriptor lists, protection look-up to enable the state update of the commands and enabling the flow of
20 the data associated with the commands to or from the host processor through the host interface;

g. command look-up tables to store the state of active commands as directed by the said command look-up engine or retrieve the stored state as directed by the said command look-up engine;

25 h. said host interface enabling the transfer of data and/or commands to or from the host processor;

i. a host data pre-fetch manager that directs the pre-fetch of the data anticipated to be required based on the commands active inside the said IP processor, to

accelerate the data transfer to the appropriate packet processors when required for processing;

- 5 j. an output data queue for transporting the retrieved data and/or commands from the host processor to the said scheduler or the said control plane processor or the said packet processors for further processing by those resources;
- k. output buffers to hold the data received from the host using the host interface for sending them to the appropriate IP processor resources, including the scheduler or the packet processors or the control plane processor or the session cache and memory controller;
- 10 l. an output queue controller that controls the flow of the received host data to the said output buffers and the said output queues working with the host data prefetch manager and/or the command scheduler, state controller and sequencer;
- m. an input data queue for receiving incoming data extracted by the said packet processor or the control plane processor to be directed to the host processor;
- 15 n. an RDMA engine to perform the RDMA tasks required by those incoming or outgoing RDMA commands or RDMA data, comprising means for recording, retrieving and comparing the region ID, protection keys, performing address translation, and retrieving or depositing data from or to the data buffers in the host memory and further comprising means for providing instructions for the next set of actions or processes to be
- 20 active for the RDMA commands;
- o. an RDMA look-up table that holds RDMA information, including state per connection or command, used by said RDMA engine to process RDMA commands;
- p. a Target/Initiator Table which is used to record target/initiator information for said data or commands, including IP address, the port to use to reach said IP
- 25 address and connection or connections to the target/initiator used by the said command scheduler, state controller and sequencer; or
- q. a combination of any of the foregoing.

312. The process of scheduling and sequencing commands and tasks to a scheduler, control plane processor, session cache and memory controller, packet processors and other execution resources of an IP hardware processor comprising:

retrieving the commands from a command queue and interpreting the
5 command; and

retrieving the command state by retrieving the command state of execution from a command look-up engine; or storing the command initial state to the command look-up engine and command look-up tables for new command; or storing the command state to the command look-up engine and command look-up table for an active command; and

10 transmitting said commands to said scheduler, or control plane processor or session cache and memory controller, or packet processors and managing the execution of said commands through their states until command execution is completed;

313. The process of claim 312 of scheduling and sequencing commands and tasks to a host processor, comprising:

15 receiving the commands from a payload extracted from an incoming IP network packet from one of the said execution resources of the said hardware IP processor and interpreting the command; and

retrieving the command state by retrieving the command state of execution from a command look-up engine; or storing the command initial state to the command look-up engine and command look-up tables for a new command; or storing the command state to the command look-up engine and command look-up table for an active command; and
20

scheduling the received command to a new requests queue or a completion queue to be sent to the host processor by the host interface process.

314. The process of claim 312 where in the said command queue is a new
25 command queue that receives new commands from the host processor to be processed by the said IP processor;

315. The process of claim 312 where in the said command queue is an active commands queue that holds the commands that are being executed during their using the

resources of the said packet processor, including issuing communications over the IP network with the peer and processing received responses from the IP network.

316. The process of claim 312 further comprising interpreting the said command and identifying the command to be an RDMA command and managing the command
5 through the execution process by a RDMA engine which performs the process of RDMA actions;

317. An RDMA process comprising: identifying the RDMA command and the connection that it is associated with; retrieving the state of the RDMA command; selecting the next step for handling the RDMA command based on the current RDMA state and
10 retrieving associated data from a host data buffer or queuing said associated data for depositing to the RDMA buffer of a host processor identified by the said RDMA command; and updating the state of the RDMA process in an RDMA look-up table for use in handling the next command associated with this connection for storage or non-storage data transfer.

318. The process of claim 316 further comprising an RDMA engine requesting the
15 command scheduler process of delivering the received RDMA command to the host through new commands for establishing a new RDMA connection with a peer.

319. The process of claim 316 further comprising an RDMA engine requesting the command scheduler process to retrieve data from a host RDMA buffer associated with the RDMA command or depositing the data associated with the RDMA command to the host
20 RDMA buffer associated with the RDMA command.

320. The process of claim 317 further comprising retrieving the region ID, access or protection keys associated with the connection that the RDMA command belongs to and performing the access or protection key verification with that received with the said RDMA command and if the access control is valid, performing the address translation from the
25 region ID to the virtual or physical address of the host RDMA buffer to be used in the data transfer.

321. The process of claim 320 further comprising retrieving the data associated with said RDMA command having valid access control from the said host RDMA buffer for sending the data to the resources of the said IP processor for transporting them to their
30 destination over the IP network.

322. The process of claim 320 further comprising receiving the data associated with the RDMA command having valid access control from the various resources of the IP packet processor extracted from the received IP packet and sending the data to the host RDMA buffers after performing address translation to point to the RDMA buffer;

5 323. The process of claim 312 further comprising a host data pre-fetch manager retrieving the data associated with an outgoing command that is to be fetched when requested by the said resources of the IP processor, or therefore, from the host buffers using the host controller interface.

10 324. The process of claim 323 further comprising the said host data pre-fetch manager working with an output queue controller to direct the data received from the host buffers to the output data queue.

15 325. The process of claim 312 further comprising retrieving or storing the target or the initiator information from or to the target/initiator table to identify the associated media interface output port and the IP address and media interface port for the target/initiator involved in this command.

326. The process of claim 313 further comprising retrieving or storing the target or the initiator information from or to the target/initiator table to identify the associated media interface output port and the IP address and media interface port for the target/initiator involved in this command.

20 327. The process of claim 312 further comprising receiving the data associated with the command from the various resources of the said IP packet processor where data is extracted from the received IP packet and sending the data to the host buffers.

25 328. The process of claim 313 further comprising receiving the data associated with the command from the various resources of the said IP packet processor where data is extracted from the received IP packet and sending the data to the host buffers.

329. For use in a hardware implemented IP network application processor, including execution resources, a host interface controller comprising a host bus interface controller to control the physical protocol for transporting data to and from a host bus; a host transaction controller and interrupt controller for controlling and directing transactions on a

host bus used to perform data transfers over the host bus; a DMA engine used to perform direct memory access of the data involved in the transfers directly to or from a host memory without substantial host processor intervention; a host command unit used to decode the command received from the drivers or applications on a host processor involved in the data transfer over an IP network including commands to setup or retrieve various configuration, control or data register resources of the said IP network application processor; and a host input queue interface providing the data received from the host to be provided to the resources of the said IP network application processor; a host output queue interface providing the data received from the resources of the processor for depositing them into the host memory; and a host command queue interface to provide the commands to the resources of the IP network application processor.

330. The combination of claim 329 where the host bus is a PCI, PCI-X, 3GIO, PCI-Express, AMBA Bus, Rapid IO or HyperTransport, or a combination of any of the foregoing.

331. The combination of claim 329 where the host bus is a host bus for switching system fabric interfaces comprising CSIX bus, XPF, XAUI interface, GMII, MII, XGMII interface, SPI interface, SPI-4 or other SPI derivative interface, PCI, PCI-Express, Infiniband, Fibre channel, RapidIO or Hypertransport or other proprietary, standard or non-standard buses.

332. For use in a hardware implemented IP network application processor, a security engine, comprising: at least one of an authentication engine for providing message digest and message authentication capabilities; an encryption/decryption engine that provides various encryption/decryption algorithm implementations to encrypt outgoing data or decrypt incoming data with the appropriate algorithm; a sequencer that sequences incoming packets through at least one of authentication and encryption/decryption engines and is used to fetch the appropriate security context for the packet being processed; a security context memory used to hold a security association database for various connections that require security operations; a coprocessor interface and queue manager used to interface a security engine and said security context with an offchip security processor and/or security context memory; one or more clear packet input queues to receive packets that need encryption and/or message authentication; a secure packet output queue used to transfer the packets that have gone through security processing by the security engine on their way out to the IP network from the said IP network application processor; a

secure packet input queue which receives the incoming IP network packets that are classified as secure packets and need security processing before further processing inside the IP network application processor; or a clear packet output queue used to transfer the incoming IP network packets that have been processed by the security engine, or a
5 combination of any of the foregoing.

333. The combination of claim 332 wherein said incoming packets further comprises a tag prepended to the packet indicating the type of operation or service required by the packet as identified by the classifier or the execution resources of the said IP network application processor.

10 334. The combination of claim 332 wherein said authentication engine further comprises one or more capabilities to perform message authentication or message digest algorithms as security service requested by the tag attached to the packet.

335. The combination of claim 334 wherein the message authentication algorithms are secure hash algorithms and the message digest algorithms are MD4, MD5 or follow-ons
15 thereof.

336. The combination of claim 332 wherein said encryption/decryption engine further comprises one or more capabilities to perform message encryption or decryption using one or more encryption decryption algorithms as security service requested by the said tag attached to the packet.

20 337. The combination of claim 332 wherein the encryption algorithms are DES, 3DES, or AES, or follow-ons thereof.

338. For use in a hardware implemented IP network application processor comprising one or more packet processor engines, a session controller or connection manager comprising a global session cache and memory complex for caching, storing and
25 retrieving session database entries for the connections being processed by the said IP network application processor; a control plane processor to create and teardown session entries to be held in the session cache and memory; a local session database cache to hold the active session information inside the packet processor engines of the said IP network application processor; a session database lookup engine inside one or more of the packet
30 processor engines to retrieve session database entries from the global session cache and

memory; a session manager that is local to a packet processor used to retrieve session entries from local session cache; and a global session data base look-up engine inside the session cache and memory complex to store, search and retrieve specific session entries to serve the specific sessions from a session memory.

5 339. The combination of claim 338 wherein the at least one packet processor engine comprises a RISC-style processor or a RISC processor with network oriented instruction set enhancement or a TCP/IP processor engine or an IP Storage processor engine or a general purpose processor or a combination of any of these processors.

10 340. The combination of claim 338 wherein the control plane processor comprises a RISC-style processor, a system controller, a micro-controller or a state machine implemented processor or a general purpose processor or a combination of any of these processors.

15 341. The combination of claim 340 wherein the state machine implemented processor includes the ability to perform functions that are full session initiation, session entry creation, session teardown, session entry removal, error handling, session relocation or session entry search and retrieval from the session memory, or a combination of any of the foregoing.

20 342. The combination of claim 338 wherein the session database cache includes the session information and fields from TCP/IP sessions, IP Storage sessions, RDMA sessions or other network oriented connections between the network source and destination systems.

25 343. The combination of claim 338 wherein the session cache and memory complex includes a cache array to hold session entries and pointers to session fields held off-chip, comprising a tag array includes tag address fields used to match a hash index or a session entry look-up index address to find an associated cache entry from a memory array; tag match logic to compare the index address matching with the tag address fields; memory banks to hold the session entry fields including pointers to fields held in off-chip memory; memory row and column decoders to decode the address of the session entry to be retrieved or stored; data input/output ports and buffers to hold session entry data to be
30 written or read from the session memory arrays; a session look-up engine to search session entries based on a session index to read or write said entries and/or their specific fields; and

an external memory controller to store and retrieve session entries and/or other IP packet processor data to/from off chip memories.

344. The session index of claim 343 comprising an index derived using a hashing or an algebraic or geometric hash or binary search algorithm.

5 345. The combination of claim 343 wherein said memory on a chip not included in said IP network application processor is RAM, DRAM, SDRAM, DDR SDRAM, RDRAM, FCRAM, FLASH, ROM, EPROM, EEPROM, QDR SRAM QDR DRAM or other derivatives of static or dynamic random access, or a combination of any of the foregoing.

10 346. The session cache and memory complex of claim 343 further comprising a debug and self test component to debug the cache tag arrays and memory arrays during silicon testing or at reset or power-up; and a direct access test structure used to test the integrity of the cache array during manufacturing or at reset or power-up or other system validation environments.

15 347. A switching system comprising a plurality of line cards coupled to a switching fabric, said line cards including a processor capable of execution transport layer RDMA functions for processing Internet data packets in one or more sessions, said processor including a session memory for storing frequently or recently used session information for a plurality of sessions.

20 348. The switching system of claim 347 wherein said processor on at least one line card includes a host interface functioning as a fabric interface and said switching fabric is coupled to said fabric interface through a fabric controller.

349. The switching system of claim 347 wherein said fabric controller functions as an interface to said switching fabric, as a traffic manager as a traffic shaper or a combination of any of the foregoing.

25 350. An IP processor having transport layer RDMA capability and comprising an IP network application processor core or an IP Storage network application processor core for enabling TCP over IP networks, said processor core comprising:

a. an RDMA mechanism for performing RDMA data transfer

- b. at least one packet processor for processing packets;
- c. a session memory for storing session information;
- d. at least one memory controller for controlling memory accesses;
- e. a media interface for coupling to at least one network; and
- 5 f. a host interface for coupling to at least one host or a fabric interface for coupling to a fabric.

351. The IP processor of claim 350 further comprising at least one of:

- a. an IP Storage session memory for storing IP Storage session information;
- 10 b. a classification processor for classifying IP packets;
- c. a flow controller for controlling data flow;
- d. a policy processor for applying policies;
- e. a security processor for performing security operations;
- f. at least one packet memory for storing packets
- 15 g. a controller for control plane processing; and
- h. a combination of any of the foregoing.

352. The IP processor of claim 350 wherein any combination of said recited elements a through f or parts thereof are implemented in a single element.

20 353. The IP processor of claim 351 wherein any combination of said recited elements a through h or parts thereof are implemented in a single element.

354. The processor of claim 2 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS,

DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

355. The processor of claim 3 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP,
5 infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

356. The processor of claim 4 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP,
10 infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

357. The processor of claim 22 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP,
15 infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

358. The processor of claim 29 wherein said hardware processor is programmable and operates on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP,
20 infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivatives, SGML, or HTML format, or a combination of any of the foregoing.

359. The combination of claim 32 wherein said processor is included within a microcontroller, a processor or a chipset of at least one of said apparatus.

25 360. The processor of claim 42 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using a iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivatives, SGML, or HTML format, or a combination of any of the foregoing.

361. The combination of claim 47 wherein said processor itself includes a processor for performing deep packet inspection and classification.

362. The combination of claim 48 wherein said processor itself includes a processor for performing deep packet inspection and classification.

5 363. The processor of claim 53 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

10 364. The processor of claim 54 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

15 365. The combination of claim 58, said processor having a security engine or a classification engine, or a combination of said engines, said engines being on separate chips of said chip set.

366. The combination of claim 61, said queues and said controller being implemented on a chip in said chipset other than a chip on a mother board of said host.

20 367. A hardware processor capable of executing a transport layer RDMA protocol on an IP Network.

368. The manufacturing process of creating a hardware processor capable of executing a transport layer RDMA protocol on an IP Network.

25 369. The manufacturing process of creating on a hardware processor an RDMA mechanism capable of performing a transport layer RDMA protocol on an IP Network.

370. The combination of a peer system connected to a host system, each of said peer system and said host system including at least one hardware processor capable of executing a transport layer RDMA protocol on an IP Network, said combination capable of

performing (a) RDMA transfer from the peer system to the host system, (b) RDMA transfer from the host system to the peer system, and (c) RDMA transfer from the peer system to the host system and from the host system to the peer system concurrently.

371. A hardware processor capable of executing a transport layer RDMA protocol
5 on an IP Network for Internet Protocol data transfers, and including at least one packet processor having an internal memory containing as database entries, frequently or recently used IP session information for processing said data transfers.

372. The combination of claim 371 further comprising a global memory located on
10 or off said hardware processor and accessible by said at least one packet processor, said global memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said internal memory.

373. The combination of claim 372 further comprising an external memory located
15 on or off said hardware processor and coupled to said global memory, said external memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said external memory being less frequently or less recently used than the IP session information contained in said global memory.

374. A hardware processor having at least one packet processor engine and
20 capable of executing a transport layer RDMA protocol over an IP Network for Internet Protocol data transfers, and including:

- a. a first memory internal to said at least one packet processor and containing as database entries frequently or recently used IP session information for processing said data transfers;
- 25 b. a global memory located on or off said hardware processor, said global memory accessible by said at least one packet processor and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said first memory; and

c. a third memory located on or off the hardware processor, said third memory coupled to said global memory and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said third memory being less frequently or less recently used than the IP session information
5 contained in said global memory.

375. A hardware processor capable of executing a transport layer RDMA protocol on an IP Network for Internet data transfer, and including at least one packet processor engine for processing said data transfers.

376. The combination of claim 375 further comprising a global memory located on
10 or off said hardware processor, said global memory coupled to said at least one packet processor for use as a memory by said packet processor engine, said global memory containing as database entries IP session information for processing said data transfers, said IP session information in said global memory being frequently or recently used IP session information for said data transfer.

377. The combination of claim 376 further comprising an additional memory
15 located on or off said hardware processor and coupled to said global memory, said additional memory containing as database entries, IP session information, said IP session information contained in said additional memory being less frequently or less recently used than the IP session information contained in said global memory.

20 378. A hardware processor having at least one packet processor engine and capable of executing a transport layer RDMA protocol over an IP Network for Internet Protocol data transfers, and including:

a. a global memory located on or off said hardware processor, said
global memory coupled to said at least one packet processor engine for use as a memory of
25 said at least one packet processor, said global memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being frequently or recently used IP session information for said data transfers; and

b. an additional memory located on or off the hardware processor and
30 coupled to said global memory, said additional memory containing as database entries, IP

session information being less frequently or less recently used than the IP session information contained in said global memory.

379. A data processing apparatus capable of executing a transport layer RDMA protocol over IP networks for Internet Protocol data transfers during Internet sessions,
5 including:

a. at least one memory in said data processing apparatus for containing as database entries frequently or recently used IP session information for processing said Internet Protocol data transfers;

10 b. a global memory located on or off said data processing apparatus, said global memory coupled to said at least one internal memory and containing as database entries, IP session information for processing said Internet Protocol data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said at least one internal memory.

380. The data processing apparatus of claim 379 coupled to an additional memory
15 located on or off said data processing apparatus, said additional memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said additional memory being less frequently or less recently used than the IP session information contained in said global memory.

381. A data processing apparatus for processing Internet Protocol data transfers
20 during Internet sessions, including:

a. at least one internal memory on said data processing apparatus containing as database entries, frequently or recently used IP session information for processing said data transfers;

25 b. a global memory on said data processing apparatus or off said data processing apparatus, said global memory coupled to said at least one internal memory and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently used or less recently used than the IP session information contained in said at least one internal memory; and

c. an additional memory located on or off said data processing apparatus, said additional memory containing as database entries, IP session information for processing said data transfers said IP session information contained said additional memory being less frequently or less recently used than the IP session information contained in said global memory.

382. For use in a processor capable of executing a transport layer RDMA protocol on an IP Network for Internet Protocol data transfers in IP sessions, a session memory containing as database entries, frequently or recently used IP session information for processing said data transfers.

383. The session memory of claim 382 further comprising a global memory accessible by said session memory, said global memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said session memory.

384. The combination of claim 383 further comprising a third memory accessible by said global memory, said third memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said third memory being less frequently or less recently used than the IP session information contained in said global memory.

385. For use in a processor capable of executing a transport layer RDMA protocol over an IP Network for Internet Protocol data transfers in IP sessions a memory system comprising:

- a. a session memory containing as database entries frequently or recently used IP session information for processing said data transfers;
- b. a global memory accessible by said session memory and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said session memory; and

c. a third memory accessible by global memory and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said third memory being less frequently or less recently used than the IP session information contained in said global memory.

5 386. A server that is a blade server, thin server, appliance server, unix server, linux server, Windows or Windows derivative server, clustered server, database server, grid computing server, VOIP server, wireless gateway server, security server, file server, network attached server, media server, streaming media server or game server, or a combination of any of the foregoing, said server including a chipset containing a hardware processor
10 providing a transport layer remote direct memory access capability over TCP, SCTP, UDC or other session oriented protocol on a network.

 387. A server that is a blade server, thin server, appliance server, unix server, linux server, Windows or Windows derivative server, clustered server, database server, grid computing server, VOIP server, wireless gateway server, security server, file server, network
15 attached server, media server, streaming media server or game server, or a combination of any of the foregoing, said server including a chipset containing a hardware processor providing remote direct memory access capability over a protocol other than TCP, SCTP or UDP.

 388. The server of claim 386 wherein said processor is programmable and
20 operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

 389. The server of claim 387 wherein said processor is programmable and
25 operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

 390. The server of claim 388 wherein said processor has certain of its functions
30 implemented in hardware and certain of its functions implemented in software.

391. The server of claim 389 wherein said processor has certain of its functions implemented in hardware and certain of its functions implemented in software.

392. The server of claim 390, said processor included as a companion processor on said server chipset.

5 393. The server of claim 391, said processor included as a companion processor on said server chipset.

394. A storage controller for controlling storage and retrieval to and from a storage area network, of data transmitted over IP networks, said storage controller including a hardware processor providing a transport layer remote direct memory access capability for enabling storage using TCP, SCTP or UDP over IP.

10

395. The storage controller of claim 394 wherein said hardware processor is included as a companion processor on a chipset of said storage controller.

396. The storage controller of claim 394 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

15

397. The storage controller of claim 395 wherein said processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

20

398. The combination of claim 394 wherein said processor (1) is embedded on a chipset on the storage controller's motherboard, or (2) includes the function of data packet security, or (3) includes the function of data packet scheduling, or (3) includes the function of data packet classification.

25

399. The combination of claim 395 wherein said processor (1) is embedded on a chipset on the storage controller's motherboard, or (2) includes the function of data packet

security, or (3) includes the function of data packet scheduling, or (3) includes the function of data packet classification.

400. The combination of claim 394 wherein said processor provides IP network storage capability for said storage controller to operate in an IP based storage area network.

5 401. The combination of claim 395 wherein said processor provides IP network storage capability for said storage controller to operate in an IP based storage area network.

402. The combination of claim 394 wherein said storage controller includes a blade controller having an interface to said storage area network.

10 403. The combination of claim 397 wherein said storage controller includes a blade controller having an interface to said storage area network and said chipset is in said interface.

404. The combination of claim 394 further comprising at least one storage array, said storage controller providing access said at least one storage array and controlling the storage function in said at least one storage array.

15 405. The combination of claim 395 further comprising at least one storage array, said storage controller providing access said at least one storage array and controlling the storage function in said at least one storage array.

20 406. The combination of claim 386 wherein said server includes a host adapter card and said processor is embedded in said host adapter card for providing high-speed TCP/IP networking capability.

407. The combination of claim 387 wherein said server includes a host adapter card and said processor is embedded in said host adapter card for providing high-speed networking capability.

25 408. The combination of claim 406 wherein said host adapter card is capable of accessing network storage to transmit data to said storage over the Internet.

409. The combination of claim 407 wherein said host adapter card is capable of accessing network storage to transmit data to said storage over the Internet.

410. The combination of claim 406 wherein said adapter card is used as a blade in a scalable blade server.

411. The combination of claim 407 wherein said host adapter card is used as a blade in a scalable blade server.

5 412. The combination of claim 406 wherein said host adapter card is in the front end of a storage array front end.

413. The combination of claim 407 wherein said host adapter card is in the front end of a storage array front end.

10 414. A host processor for processing data packets received over the Internet, said host processor including a hardware processor providing a transport layer remote direct memory access capability for enabling storage using TCP, SCTP or UDP or other session oriented protocol, over IP networks, said hardware processor providing offloading capability.

15 415. The combination of claim 414 wherein said hardware processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, Infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a combination of any of the foregoing.

20 416. The combination of claim 414 wherein said host processor is a high end server, workstation, personal computers capable of interfacing with high speed networks, wireless LAN, hand held wireless telecommunication device, router, switch, gateway, blade server, thin server, media server, streaming media server, appliance server, Unix server, Linux server, Windows or Windows derivative server, AIX server, clustered server, database server, grid computing server, VOIP server, wireless gateway server, security server, file server, network attached storage server, game server or a combination of any of the
25 foregoing.

417. The combination of claim 415 wherein said host processor is a high end server, workstation, personal computers capable of interfacing with high speed networks, wireless LAN, hand held wireless telecommunication device, router, switch and gateway, blade server, thin server, media server, streaming media server, appliance server, Unix

server, Linux server, Windows or Windows derivative server, AIX server, clustered server, database server, grid computing server, VOIP server, wireless gateway server, security server, file server, network attached storage server, game server or a combination of any of the foregoing.

5 418. The combination of claim 416 wherein at least one member of said group of recited apparati is a low power apparatus.

 419. The combination of claim 417 wherein at least one member of said group of recited apparati is a low power apparatus.

 420. An IP storage area network switching system line card having embedded
10 therein a hardware processor providing remote direct memory access capability for enabling high-speed storage using TCP, SCTP or UDP over IP networks, said processor being programmable and operating on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivative, SGML, or HTML format or a
15 combination of any of the foregoing.

 421. A gateway controller of a storage area network, said gateway controller including a chipset having embedded therein a hardware processor providing a transport layer remote direct memory access capability for enabling high-speed storage using TCP, SCTP or UDP over IP networks.

20 422. The combination of claim 421 wherein said hardware processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivatives, SGML, or HTML format or a combination of any of the foregoing.

25 423. A storage area network management appliance including a chipset having embedded therein a hardware processor providing a transport layer remote direct memory access capability for enabling storage of traffic using TCP, SCTP or UDP over IP networks, said hardware processor enabling said appliance to transport TCP/IP packets in-band to said traffic or out of band to said traffic.

424. A cluster of servers, each server including at least one chipset, at least one of said chipsets in said cluster having embedded therein a hardware processor providing remote direct memory access capability for enabling high-speed storage using TCP, SCTP or UDP over IP networks.

5 425. The combination of claim 424 wherein said hardware processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivatives, SGML, or HTML format or a combination of any of the foregoing.

10 426. A chip set having embedded therein a hardware processor providing transport layer remote direct memory access capability for enabling high-speed storage using TCP, SCTP or UDP over IP networks.

15 427. The chip set of claim 426 wherein said hardware processor is programmable and operates on data packets transmitted, encapsulated or encoded using an iSCSI, iFCP, infiniband, SATA, SAS, IP, ICMP, IPSEC, DES, 3DES, AES, FC, SCSI, FCIP, NFS, CIFS, DAFS, HTTP, XML, XML derivatives, SGML, or HTML format or a combination of any of the foregoing.

 428. The combination of claim 426 wherein said hardware processor has certain of its functions implemented in hardware and certain of its functions implemented in software.

20 429. The combination of claim 426, said hardware processor having a security engine or a classification engine or a combination of either.

 430. A host processor running an application, said host processor including a hardware processor providing transport layer remote direct memory access (RDMA) capability, said hardware processor implemented for enabling high-speed storage using
25 TCP, SCTP or UDP over IP networks, said hardware processor including:

 a. registration circuitry for allowing said application to register a memory region of said host processor with said hardware processor for RDMA access;

 b. communication circuitry for allowing the exporting of said registered memory region to at least one peer hardware processor having RDMA capability and for

allowing the informing of said peer of said host processor's desire to allow said peer to read data from or write data to said registered memory region; and

c. RDMA circuitry for allowing information transfer to and from said registered region of memory without substantial host processor intervention.

5 431. A host processor running an application, said host processor including a hardware processor providing transport layer remote direct memory access (RDMA) capability, said hardware processor implemented on an integrated circuit chip for enabling high-speed storage using TCP, SCTP or UDP over IP networks, the process of performing RDMA for said application including the steps of:

10 a. said application registering a region of memory of said host processor for RDMA;

 b. said hardware processor making said region of memory available to a peer processor having remote data transfer access capability without substantial intervention by said host processor in said data transfer;

15 c. said hardware processor communicating to said peer processor said host processor's desire to allow said peer processor to read data from or write data to said region of memory; and

 d. said hardware processor enabling information transfer from or to said registered region of memory without said host processor's substantial intervention in said
20 information transfer.

 432. A host processor having a SCSI command larger and an iSCSI driver, said host processor including a hardware processor including a transport layer RDMA capability for providing TCP/IP operations over a network for data packets from or to an initiator and providing said packets to or from said target, said operations requested by a host processor,
25 said hardware processor comprising:

 a. an RDMA mechanism;

 b. a command scheduler for scheduling commands from the command layer of said host processor for operation in said hardware processor;

142

- c. first command queues for queuing commands from said host processor for existing sessions;
- d. second command queues for queuing commands from said host processor for sessions that do not currently exist;
- 5 e. a database for recording the state of the session on which said command is transported, said database also for recording progress of RDMA for those of said commands that use RDMA;
- f. a communication path between said hardware processor and said SCSI layer of said host processor for communicating status of command execution to said
10 SCSI layer for processing; and
- g. at least one transmit/receive engine and at least one command engine coupled together, said engines working together to interpret commands and perform appropriate operations for performing RDMA for retrieving data from or transmitting data to said host processor and for updating said state of said session.
- 15 433. The combination of claim 432 wherein said at least one transmit/receive engine is implemented as a separate transmit engine and a separate receive engine.
434. The combination of claim 432 wherein said first command queues are located partly in memory on said hardware processor and partly in memory off said hardware processor.
- 20 435. The combination of claim 432 wherein said second command queues are located partly in memory on said hardware processor and partly in memory off said hardware processor.
436. The combination of claim 434 wherein said memory off said hardware processor is memory in said host processor and memory on a chip not included in said host
25 processor.
437. The combination of claim 435 wherein said memory off said hardware processor is memory in said host processor and memory on a chip not included in said host processor.

438. The combination of claim 436 wherein said memory on a chip not included in said host processor is RAM, DRAM, SDRAM, DDR SDRAM, RDRAM, FCRAM, FLASH, ROM, EPROM, EEPROM, QDR SRAM QDR DRAM and other derivatives of static or dynamic random access memory, or a combination of any of the foregoing.

5 439. The combination of claim 436 wherein said memory on a chip not included in said host processor is located on a companion chip to said hardware processor.

 440. The combination of claim 437 wherein said memory on a chip not included in said host processor is RAM, DRAM, SDRAM, DDR SDRAM, RDRAM, FCRAM, FLASH, ROM, EPROM, EEPROM, QDR SRAM QDR DRAM and other derivatives of static or
10 dynamic random access memory, or a combination of any of the foregoing.

 441. The combination of claim 437 wherein said memory on a chip not included in said host processor is located on a companion chip to said hardware processor.

 442. A host processor having a SCSI command layer and an iSCSI driver, said host processor capable of being coupled to a hardware implemented iSCSI controller
15 useable in high speed storage over IP, said controller for transporting received iSCSI commands and PDUs, said controller having access to a data base for keeping track of data processing operations, said database being in memory on said controller, or in memory partly on said controller and partly in a computing apparatus other than said controller, said controller having a transmit and a receive path for data flow, said controller comprising:

20 a. a command scheduler for scheduling processing of commands, said scheduler coupled to said SCSI command layer and to said iSCSI driver;

 b. a receive path for data flow of received data and a transmit path for data flow of transmitted data;

 c. at least one transmit engine for transmitting iSCSI PDUs;

25 d. at least one transmit command engine for interpreting said PDUs and performing operations including retrieving information from said host processor using remote direct memory access and keeping command flow information in said database updated as said retrieving progresses;

e. at least one receive command engine; and

f. at least one receive engine for interpreting received commands into requests for at least one of said at least one receive command engine.

5 443. A storage array controller having a chipset for controlling a storage array, said controller including a hardware processor capable of providing remote direct memory access capability for enabling storage using TCP, SCTP or UDP over IP networks, said processor controlling storage and retrieval, to and from said storage array, of data transmitted over IP networks, said processor included as a companion processor on said chipset

10 444. The combination of claim 443 wherein said processor provides IP network storage capability for said storage array controller to operate in an IP based storage area network.

15 445. In a hardware processor capable of executing a transport layer RDMA protocol for transporting data packets in TCP/IP or other session oriented protocol sessions or flows over an IP network, a scheduler for scheduling said packets to execution resources of the hardware processor, said scheduler comprising:

a. a resource allocation table for storing

i. the identification of at least some of said execution resources,

ii. the identification of the session which at least one of said resources is executing, and

20 iii. the processing state of said resources, and

b. a state controller and sequencer coupled to said resource allocation table and to said execution resources, said state controller and sequencer scheduling packets to be processed on a specific session to the execution resource executing said specific session, said scheduling based on at least the execution state of said execution resource.

25

446. The scheduler of claim 445 wherein said resource allocation table also stores the TCP/IP session on which at least some of said packets are to be processed.

447. The scheduler of claim 445 including a packet memory for storing a plurality of packets, said packet memory coupled to said state controller and sequencer.

448. In a data processing apparatus capable of executing a transport layer RDMA protocol for transporting data packets in TCP/IP or other session oriented protocol sessions
5 or flows over an IP network, a scheduler for scheduling said packets to execution resources of the hardware processor, said scheduler comprising:

- a. a resource allocation table for storing
 - i. the identification of at least some of said execution resources,
 - ii. the identification of the TCP/IP session which at least one of
10 said resources is executing,
 - iii. the processing state of said resources, and
- b. a state controller and sequencer coupled to said resource allocation table, and to said execution resources, said state controller and sequencer scheduling packets to be processed on a specific session to the execution resource executing said
15 specific session, said scheduling based on at least the execution state of said execution resource.

449. The scheduler of claim 449 wherein said resource allocation table also stores the session on which at least some of said packets are to be processed.

450. The scheduler of claim 448 further comprising a packet memory for storing a
20 plurality of packets, said packet memory coupled to said state controller and sequencer.

451. A switching system comprising a plurality of line cards coupled to a switching fabric, said line cards including a processor for processing Internet data packets in one or more sessions, said processor including a session memory for storing frequently or recently used session information for a plurality of sessions.

25 452. The switching system of claim 451 wherein said processor on at least one line card includes a host interface functioning as a fabric interface and said switching fabric is coupled to said fabric interface through a fabric controller.

453. The switching system of claim 452 wherein said fabric controller functions as an interface to said switching fabric, as a traffic manager as a traffic shaper or a combination of any of the foregoing.

5 454. An IP processor for enabling TCP or SCTP, or UDP, or other session oriented protocols over IP networks, said IP processor comprising:

- a. at least one packet processor for processing IP packets;
- b. a session memory for storing IP session information;
- c. at least one memory controller for controlling memory accesses;
- d. at least one media interface for coupling to at least one network; and
- 10 e. a host interface for coupling to at least one host or a fabric interface for coupling to a fabric.

455. The IP processor of claim 454 further comprising at least one of:

- a. an IP Storage session memory for storing IP Storage session information;
- 15 b. a classification processor for classifying IP packets;
- c. a flow controller for controlling data flow;
- d. a policy processor for applying policies;
- e. a security processor for performing security operations;
- f. a packet memory for storing packets
- 20 g. a controller for control plane processing;
- h. a packet scheduler;
- i. a connection manager or session controller for managing sessions; or
- j. a combination of any of the foregoing.

456. A hardware processor for Internet Protocol data transfers, said processor including at least one packet processor having an internal memory containing as database entries, frequently or recently used IP session information for processing said data transfers.

5 457. The combination of claim 456 further comprising a global memory located on or off said hardware processor and accessible by said at least one packet processor, said global memory containing as database entries, session information for processing said data transfers, said session information contained in said global memory being less frequently or less recently used than the session information contained in said internal memory.

10 458. The combination of claim 457 further comprising an external memory located on or off said hardware processor and coupled to said global memory, said external memory containing as database entries, session information for processing said data transfers, said session information contained in said external memory being less frequently or less recently used than the IP session information contained in said global memory.

15 459. A hardware processor having at least one packet processor engine for Internet Protocol data transfers over an IP Network, said hardware processor including:

- a. a first memory internal to said at least one packet processor engine and containing as database entries frequently or recently used IP session information for processing said data transfers;
- 20 b. a global memory located on or off said hardware processor, said global memory accessible by said at least one packet processor engine and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said first memory; and
- 25 c. a third memory located on or off the hardware processor, said third memory coupled to said global memory and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said third memory being less frequently or less recently used than the IP session information contained in said global memory.

460. A hardware processor for Internet data transfer on an IP Network, and including at least one packet processor engine for processing said data transfers and further comprising a global memory located on or off said hardware processor, said global memory coupled to said at least one packet processor engine for use as a memory by said packet processor, said global memory containing as database entries IP session information for processing said data transfers, said IP session information in said global memory being frequently or recently used IP session information for said data transfer.

461. The combination of claim 460 further comprising an additional memory located on or off said hardware processor engine and coupled to said global memory, said additional memory containing as database entries, IP session information, said IP session information contained in said additional memory being less frequently or less recently used than the IP session information contained in said global memory.

462. A hardware processor having at least one packet processor engine for Internet Protocol data transfers over an IP Network, and including:

a. a global memory located on or off said hardware processor, said global memory coupled to said at least one packet processor engine for use as a memory of said at least one packet processor engine, said global memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being frequently or recently used IP session information for said data transfers; and

b. an additional memory located on or off the hardware processor and coupled to said global memory, said additional memory containing as database entries, IP session information being less frequently or less recently used than the IP session information contained in said global memory.

463. A data processing apparatus for Internet Protocol data transfers over IP networks during Internet sessions, including:

a. at least one memory in said data processing apparatus for containing as database entries frequently or recently used IP session information for processing said Internet Protocol data transfers;

b. a global memory located on or off said data processing apparatus, said global memory coupled to said at least one internal memory and containing as database entries, IP session information for processing said Internet Protocol data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said at least one internal memory.

464. The data processing apparatus of claim 463 coupled to an additional memory located on or off said data processing apparatus, said additional memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said additional memory being less frequently or less recently used than the IP session information contained in said global memory.

465. A data processing apparatus capable of processing Internet Protocol data transfers during Internet sessions, including:

a. at least one internal memory on said data processing apparatus containing as database entries, frequently or recently used IP session information for processing said data transfers;

b. a global memory on said data processing apparatus or off said data processing apparatus, said global memory coupled to said at least one internal memory and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently used or less recently used than the IP session information contained in said at least one internal memory; and

c. an additional memory located on or off said data processing apparatus, said additional memory containing as database entries, IP session information for processing said data transfers said IP session information contained said additional memory being less frequently or less recently used than the IP session information contained in said global memory.

466. For use in a processor for Internet Protocol data transfers on an IP Network in IP sessions, a session memory containing as database entries, frequently or recently used IP session information for processing said data transfers.

467. The session memory of claim 466 further comprising a global memory accessible by said session memory, said global memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said session memory.

468. The combination of claim 467 further comprising an third memory accessible by said global memory, said third memory containing as database entries, IP session information for processing said data transfers, said IP session information contained in said third memory being less frequently or less recently used than the IP session information contained in said global memory.

469. For use in a processor for Internet Protocol data transfers over an IP Network in IP sessions, a memory system comprising:

- a. a session memory containing as database entries frequently or recently used IP session information for processing said data transfers;
- b. a global memory accessible by said session memory and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said global memory being less frequently or less recently used than the IP session information contained in said session memory; and
- c. a third memory accessible by global memory and containing as database entries, IP session information for processing said data transfers, said IP session information contained in said third memory being less frequently or less recently used than the IP session information contained in said global memory.

470. A host processor having a SCSI command layer and an iSCSI or IP Storage driver, said host processor capable of being coupled to a hardware implemented iSCSI or IP Storage controller useable in high speed storage over IP, said controller for transporting received iSCSI or IP Storage commands and PDUs, said controller having access to a data base for keeping track of data processing operations, said database being in memory on said controller, or in memory partly on said controller and partly in a computing apparatus other than said controller, said controller having a transmit and a receive path for data flow, said controller comprising:

151

- a. a command scheduler for scheduling processing of commands, said scheduler coupled to said SCSI command layer and to said iSCSI or IP Storage driver;
 - b. a receive path for data flow of received data and a transmit path for data flow of transmitted data;
 - 5 c. at least one transmit engine for transmitting iSCSI or IP Storage PDUs;
 - d. at least one transmit command engine for interpreting said PDUs and performing operations including retrieving information from said host processor using remote direct memory access and keeping command flow information in said database updated as
10 said retrieving progresses;
 - e. at least one receive command engine; and
 - f. at least one receive engine for interpreting received commands into requests for at least one of said at least one receive command engine.
471. In a hardware processor for transporting data packets in TCP/IP or other
15 session oriented IP Protocol sessions or flows over an IP network, a scheduler for scheduling said packets to execution resources of the hardware processor, said scheduler comprising:
- a. a resource allocation table for storing
 - i. the identification of at least some of said execution resources,
 - 20 ii. the identification of the session which at least one of said resources is executing, and
 - iii. the processing state of said resources, and
 - b. a state controller and sequencer coupled to said resource allocation table and to said execution resources, said state controller and sequencer scheduling
25 packets to be processed on a specific session to the execution resource executing said specific session, said scheduling based on at least the execution state of said execution resource.

472. The scheduler of claim 471 wherein said resource allocation table also stores the session on which at least some of said packets are to be processed.

473. The scheduler of claim 471 including a packet memory for storing a plurality of packets, said packet memory coupled to said state controller and sequencer.

5 474. In a data processing apparatus for transporting data packets in TCP/IP or other session oriented IP Protocol sessions or flows over an IP network, a scheduler for scheduling said packets to execution resources of the data processing apparatus, said scheduler comprising:

- a. a resource allocation table for storing
 - 10 i. the identification of at least some of said execution resources,
 - ii. the identification of the session which at least one of said resources is executing, and
 - iii. the processing state of said resources; and
- b. a state controller and sequencer coupled to said resource allocation
- 15 table, and to said execution resources, said state controller and sequencer scheduling packets to be processed on a specific session to the execution resource executing said specific session, said scheduling based on at least the execution state of said execution resource.

20 475. The scheduler of claim 474 wherein said resource allocation table also stores the session on which at least some of said packets are to be processed.

 476. The scheduler of claim 474 further comprising a packet memory for storing a plurality of packets, said packet memory coupled to said state controller and sequencer.

25 477. A processor for processing Internet data packets in one or more sessions, said processor including a session memory for storing frequently or recently used session information for a plurality of sessions.

 478. A TCP/IP processor implemented in hardware, said processor including a session memory for storing session information for a plurality of sessions.

479. A processor for processing Internet data packets in one or more sessions, and a session memory for storing session information for a plurality of said sessions.

480. An IP storage processor implemented in hardware, said processor including a session memory for storing session information for a plurality of sessions.

5 481. The process, in a hardware implemented control plane processor or session controller coupled to a host processor or a remote peer, of creating new sessions and their corresponding session database entries responsive to new session connection requests received either from the host processor or the remote peer.

10 482. The process, in a hardware implemented control plane processor or session controller coupled to a host processor or a remote peer and including a TCP/IP hardware processor engine or an IP storage processor engine, or a combination of any of the foregoing, of tearing down or removing sessions and their corresponding session database entries responsive to session connection closure requests received either from the host processor or the remote peer or as a result of the operation by the said TCP/IP processor engine or IP Storage processor engine or a combination of any of the foregoing.

15 483. For use in a hardware implemented IP network application processor including an input queue and queue controller for accepting incoming data packets including new commands from multiple input ports and queuing them on an input packet queue for scheduling and further processing, the process comprising: accepting incoming data packets from one or more input ports and queuing them on an input packet queue; and de-queuing said packets for scheduling and further packet processing.

20 484. For use in a hardware implemented IP network application processor having one or more output queues for accepting outgoing data packets, including new commands, from one or more packet processor engines, TCP/IP processor engines or IP Storage processing engines, directing said packets on to an output port interface for sending them out to a network interface, said interface sending said packets on to the network, through one or more output packets, said process including: accepting incoming data packets from one or more packet processing engines and queuing them on said output packet queue; and de-queuing said packets for delivery to said output port based on the destination port information in said packet.

30

485. A TCP/IP processor engine for processing Internet Protocol packets and comprising at least one of each of:

- a. a checksum hardware for performing checksum operations;
- b. a data memory for storing data used in the TCP/IP processor engine;
- 5 c. an instruction memory for storing instructions used in the TCP/IP processor engine;
- d. an instruction fetch mechanism for fetching said instructions;
- e. an instruction decoder for decoding said instructions;
- f. an instruction sequencer for sequencing said instructions;
- 10 g. a session database memory for storing TCP/IP session data; or
- h. a session database memory controller for controlling said session database memory;

or a combination of any of the foregoing items a through i; and

a host interface, or a fabric interface, or bus controller, or memory controller or combination
15 of any of the foregoing for coupling to host or a fabric.

486. The TCP/IP processor engine of claim 485 further comprising at least one of:

- a. a hash engine for performing hash functions;
- b. a sequencer manager for sequencing operations;
- c. a window operations manager for performing windowing operations to
20 position packets within, and/or verify packets to be within, agreed windows;
- d. a classification tag interpreter for interpreting classification tags;
- e. a frame controller for controlling data framing;
- f. an out of order manager for handling out of order data;

- g. a register file for storing data;
- h. a TCP state manager for managing TCP session states;
- i. a CRC component for performing CRC operations;
- j. an execution resource unit or ALU for data processing;
- 5 k. a TCP session database lookup engine for accessing session entries;
- l. an SNACK engine for selective negative acknowledgment;
- m. an SACK engine for selective positive acknowledgment;
- n. a segmentation controller for controlling the segmentation of data;
- o. a timer for event timing;
- 10 p. a packet memory for storing packets; and
- q. a combination of any of the foregoing.

487. An IP storage processor engine for processing Internet Protocol packets and comprising at least one of each of:

- a. CRC hardware for performing CRC functions;
- 15 b. a data memory for storing data used in the processor engine;
- c. an instruction memory for storing instructions used in the processor engine;
- d. an instruction fetch mechanism for fetching said instructions;
- e. an instruction decoder for decoding said instructions;
- 20 f. an instruction sequencer for sequencing said instructions;
- g. an IP storage session database memory for storing IP storage session information;

156

h. an IP storage session database memory controller for controlling said IP storage session database memory;

i. a combination of any of the foregoing items a through j; and

5 j. a host interface, or a fabric interface, or bus controller, or memory controller or combination thereof for a host or to a fabric.

488. The IP storage processor engine of claim 487 further comprising at least one of:

a. hash engine for performing hash operations;

b. a sequencer manager for sequencing operations;

10 c. a window operations manager for positioning packets within, and/or verifying received packets to be within, agreed windows;

d. a classification tag interpreter for interpreting classifications tags;

e. an out of order manager for handling out of order data;

f. a register file for storing data;

15 g. a PDU storage classifier for classifying packets into various attributes;

h. an IP storage state manager for managing IP storage session states;

i. a checksum component for performing checksum operations;

j. an execution resource unit or ALU for data processing;

20 k. an IP storage session database lookup engine for accessing session entries;

l. a SNACK engine for selective negative acknowledgment;

m. a SACK engine for selective positive acknowledgment;

n. a segmentation controller for controlling the segmentation of data;

- o. a timer for event timing;
- p. a packet memory for storing packets; and
- q. a combination of any of the foregoing.

489. A hardware implemented IP network application processor for providing
5 TCP/IP operations in sessions on information packets from or to an initiator and providing
information packets to or from a target, comprising the combination of:

- a. data processing resources including at least one programmable
packet processor for processing said packets;
- b. a TCP/IP session cache and memory controller for keeping track of
10 the progress of, and memory useful in, said operations on said packets;
- c. a host interface controller capable of controlling an interface to a host
computer in an initiator or target computer system or a fabric interface controller capable of
controlling an interface to a fabric; and
- d. a media independent interface capable of controlling an interface to
15 the network media in an initiator or target.

490. The packet processor of claim 102 further comprising (1) a packet processor
engine, or (2) a TCP/IP processor engine, or (3) an IP Storage processor engine, or a
combination of any of the foregoing.

491. The packet processor of claim 106 further comprising (1) a packet processor
20 engine, or (2) a TCP/IP processor engine, or (3) an IP Storage processor engine, or a
combination of any of the foregoing.

492. A hardware processor for enabling Internet Protocol packets or their payloads
to stream from a network interface through said hardware processor to a host interface or a
fabric interface during packet processing, said hardware processor comprising:

- 25 a. at least one packet processor for processing said packets;
- b. a packet memory;

c. a scheduler coupled to said at least one packet processor and to said packet memory for scheduling said packets to said at least one packet processor;

d. a session memory for storing session information for those packets transmitted, encapsulated or encoded using a session oriented protocol; and

5 e. a session manager coupled to the foregoing elements for managing session states for those packets transmitted, encapsulated or encoded using a session oriented protocol.

493. The hardware processor of claim 492 further including a TCP processor engine for transporting the data in said packets or an IP Storage processor engine for storing
10 the data in said packets.

494. The hardware processor of claim 492 wherein any of said elements a through e, or any parts thereof, are implemented in a single apparatus.

495. The hardware processor of claim 493 wherein any of said elements a through e or any of said processor engine, or parts thereof, are implemented in a single apparatus.

15 496. The hardware processor of claim 492 further including a processor for classifying said packets, a processor for applying a policy to said packets and a security processor for performing security functions for said packets.

497. The hardware processor of claim 496 wherein any of said processors for classifying, applying a policy or performing security functions, or parts thereof, are
20 implemented in a single apparatus.

498. The hardware processor of claim 492 wherein said packet processor is a TCP/IP processor engine, or an IP Storage processor engine, or a packet processor engine with a packet oriented instruction set, or a packet processor engine with a RISC instruction set, or a combination of any of the foregoing.

25 499. A TCP/ IP processor having transport level RDMA capability for enabling TCP or other session oriented protocols, over IP networks, said processor comprising:

a. an RDMA mechanism for performing RDMA data transfer

- b. at least one TCP/IP processor engine for processing IP packets;
- c. a session memory for storing session information;
- d. at least one memory controller for controlling memory accesses;
- e. at least one media interface for coupling to at least one network; and
- 5 f. a host interface for coupling to at least one host or a fabric interface for coupling to a fabric.

500. The TCP/IP processor of claim 499 further comprising at least one of:

- a. a packet processor engine for processing packets;
- b. a classification processor for classifying IP packets;
- 10 c. a flow controller for controlling data flow;
- d. a policy processor for applying policies;
- e. a security processor for performing security operations;
- f. a controller for control plane processing;
- g. a packet scheduler;
- 15 h. a packet memory for storing packets;
- i. a connection manager or session controller for managing sessions; or
- j. a combination of any of the foregoing.

501. A TCP/ IP processor for enabling TCP, SCTP, or UDP or other session oriented protocols, over IP networks, said processor comprising:

- 20 a. at least one TCP/IP processor engine for processing IP packets;
- b. a session memory for storing session information;
- c. at least one memory controller for controlling memory accesses;

160

- d. at least one media interface for coupling to at least one network; and
- e. a host interface for coupling to at least one host or a fabric interface for coupling to a fabric.

502. The TCP/IP processor of claim 501 further comprising at least one of:

- 5 a. a packet processor engine for processing packets;
- b. a classification processor for classifying IP packets;
- c. a flow controller for controlling data flow;
- d. a policy processor for applying policies;
- e. a security processor for performing security operations;
- 10 f. a controller for control plane processing;
- g. a packet scheduler;
- h. a packet memory for storing packets;
- i. a connection manager or session controller for managing sessions; or
- j. a combination of any of the foregoing.

15 503. An IP Storage processor having RDMA capability for enabling IP Storage protocols over IP networks, said processor comprising:

- a. an RDMA mechanism for performing RDMA data transfer;
- b. at least one IP Storage processor engine for processing IP storage
packets;
- 20 c. an IP Storage session memory for storing session information;
- d. at least one memory controller for controlling memory accesses;
- e. at least one media interface for coupling to at least one network; and

f. a host interface for coupling to at least one host or at least one fabric interface for coupling to a fabric.

504. The IP Storage processor of claim 503 further comprising at least one of:

- a. a packet processor engine for processing packets;
- 5 b. a classification processor for classifying IP packets;
- c. a flow controller for controlling data flow;
- d. a policy processor for applying policies;
- e. a security processor for performing security operations;
- f. a controller for control plane processing;
- 10 g. a packet scheduler;
- h. a packet memory for storing packets; or
- i. a combination of any of the foregoing.

505. An IP Storage processor for enabling IP Storage protocols over IP networks, said processor comprising:

- 15 a. at least one IP Storage processor engine for processing IP storage packets;
- b. an IP Storage session memory for storing session information;
- c. at least one memory controller for controlling memory accesses;
- d. at least one media interface for coupling to at least one network; and
- 20 e. a host interface for coupling to at least one host or a fabric interface for coupling to a fabric.

506. The IP Storage processor of claim 505 further comprising at least one of:

- a. a packet processor engine for processing packets;
 - b. a classification processor for classifying IP packets;
 - c. a flow controller for controlling data flow;
 - d. a policy processor for applying policies;
 - 5 e. a security processor for performing security operations;
 - f. a controller for control plane processing;
 - g. a packet scheduler;
 - h. a packet memory for storing packets; or
 - i. a combination of any of the foregoing.
- 10 507. A multiprocessor system comprising at least one data processor coupled to a plurality of IP processors for interfacing said at least one data processor to said IP processors, for enabling TCP, STCP, UDP or other session oriented protocols over IP networks, said IP processor comprising:
- a. at least one packet processor for processing IP packets;
 - 15 b. a session memory for storing IP session information;
 - c. at least one memory controller for controlling memory accesses;
 - d. at least one media interface for coupling to at least one network; and
 - e. a host interface for coupling to at least one host or fabric interface for coupling to a fabric.
- 20 508. The multiprocessor system of claim 507, said IP network application processor further comprising at least one of:
- a. an IP Storage session memory for storing IP Storage session information;

- b. a classification processor for classifying IP packets;
- c. a flow controller for controlling data flow;
- d. a policy processor for applying policies;
- e. a security processor for performing security operations;
- 5 f. a packet memory for storing packets;
- g. a controller for control plane processing;
- h. a packet scheduler;
- i. a connection manager or session controller for managing sessions' or
- j. a combination of any of the foregoing.

10 509. The multiprocessor system of claim 507 wherein two or more of said plurality of IP processors are coupled to each other.

510. The multiprocessor system of claim 508 wherein two or more of said plurality of IP processors are coupled to each other.

15 511. The multiprocessor system of claim 509 wherein said two or more of said plurality of IP processors are coupled through a co-processor interface, or a host interface, or a bridge, or a combination of any of the foregoing.

512. The multiprocessor system of claim 510 wherein said two or more of said plurality of IP processors are coupled through a co-processor interface, or a host interface, or a bridge, or a combination of any of the foregoing.

20 513. The hardware processor of claim 65 wherein said processor includes at least one engine that is a firewall engine, a router or a telecommunication network acceleration engine.

25 514. A processor implemented in hardware and capable of performing transport layer RDMA functions, and including a session memory for storing session information for a plurality of sessions.

515. The process, using a TCP/IP processor capable of executing a transport layer RDMA protocol, of transferring data from a host processor to an output media interface of a host system, said process including said TCP/IP processor forming headers and forming segmentation for said data to be transferred.

5 516. The process of claim 522 wherein at least one of said headers is a header used for transporting data through a storage area network, a local area network, a metro area network, a system area network, a system area network, a wireless local area network, a personal area network, a home network or a wide area network or a combination of any of the foregoing.

10 517. A hardware processor having transport layer RDMA capability for enabling Internet Protocol packets or their payloads to stream from a network interface through said hardware processor to a host interface or a fabric interface during packet processing, said hardware processor comprising:

- a. an RDMA mechanism for performing RDMA data transfer;
- 15 b. at least one packet processor for processing said packets;
- c. a packet memory;
- d. a scheduler coupled to said at least one packet processor and to said packet memory for scheduling said packets to said at least one packet processor;
- e. a session memory for storing session information for those packets
- 20 transmitted, encapsulated or encoded using a session oriented protocol; and
- f. a session manager coupled to the foregoing elements for managing session states for those packets transmitted, encapsulated or encoded using a session oriented protocol.

518. The hardware processor of claim 517 further including a TCP processor

25 engine for transporting the data in said packets or an IP Storage processor engine for storing the data in said packets.

519. The hardware processor of claim 517 wherein any of said elements a through e, or any parts thereof, are implemented in a single apparatus.

520. The hardware processor of claim 518 wherein any of said elements a through e or any of said processors, or parts thereof, are implemented in a single apparatus.

521. The hardware processor of claim 517 further including a processor for classifying said packets, a processor for applying a policy to said packets and a security
5 processor for performing security functions for said packets.

522. The hardware processor of claim 521 wherein any of said processors for classifying, applying a policy or performing security functions, or parts thereof, are implemented in a single apparatus.

523. The hardware processor of claim 517 wherein said packet processor is a
10 TCP/IP processor engine, or an IP Storage processor engine, or a packet processor engine with a packet oriented instruction set, or a packet processor engine with a RISC instruction set, or a combination of any of the foregoing.

524. An IP processor having RDMA capability for enabling TCP or other session oriented protocols over IP networks, said processor comprising:

- 15 a. an RDMA mechanism for performing RDMA data transfer
- b. at least one packet processor for processing IP packets;
- c. a session memory for storing IP session information;
- d. at least one memory controller for controlling memory accesses;
- e. at least one media interface for coupling to at least one network; and
- 20 f. a host interface for coupling to at least one host or a fabric interface for coupling to a fabric;
- i. wherein said processor operates in multiple stages, including one or more of the stages of (1) receiving incoming IP packets; (2) providing security for said incoming IP packets if needed; (3) classifying said incoming IP packets; (4) scheduling IP
25 packets for processing; (5) executing data processing operations on IP Packets; (6) providing direct memory access for transferring data/packets to or from the memory of a system external to said processor; (7) executing protocol processing operations on data or

commands forming IP packets; and (8) providing processing security for outgoing IP packets if needed, or (9) transmitting IP packets onto a network; or any combination of any of the foregoing; and

- (a) each of said stages is capable of operating on different
5 IP packets concurrently.

525. The IP processor of claim 524 wherein each stage of said IP processor may take a different length of time to perform its function than one or more of the other stages of said IP processor.

526. An IP processor for enabling TCP or other session oriented protocols over IP
10 networks, said processor comprising:

- a. at least one packet processor for processing IP packets;
 - b. a session memory for storing IP session information;
 - c. at least one memory controller for controlling memory accesses;
 - d. at least one media interface for coupling to at least one network; and
 - 15 e. a host interface for coupling to at least one host or a fabric interface for coupling to a fabric;
- i. wherein said processor operates in multiple stages, including one or more of the stages of (1) receiving incoming IP packets; (2) providing security for said incoming IP packets if needed; (3) classifying said incoming IP packets; (4) scheduling IP
20 packets for processing; (5) executing data processing operations on IP Packets; (6) providing direct memory access for transferring data/packets to or from the memory of a system external to said processor; (7) executing protocol processing operations on data or commands forming IP packets; and (8) providing processing security for outgoing IP packets if needed, or (9) transmitting IP packets onto a network; or any combination of any of the
25 foregoing;
 - ii. each of said stages is capable of operating on different IP packets concurrently.

527. The IP processor of claim 526 wherein each stage of said IP processor may take a different length of time to perform its function than one or more of the other stages of said IP processor.

528. The combination of claim 168 wherein said data transfer is accomplished
5 using a TCP, SCTP or UDP or other session oriented protocol, or a combination of any of the foregoing.

529. The combination of claim 168 wherein said data transfer is accomplished using a protocol selected from the group of protocols other than TCP, SCTP or UDP.

530. The IP Storage or iSCSI stack of claim 86 further comprising memory for
10 storing a database to maintain various information regarding said active sessions or connections and IP Storage or iSCSI state information for each of the sessions or connections.

531. The IP Storage or iSCSI stack of claim 86 further comprising an interface that includes circuitry for interfacing to at least one layer that is a wired or wireless Ethernet, MII,
15 GMII, XGMII, XPF, XAUI, TBI, SONET, DSL, POS, POS-PHY, SPI Interface, SPI-4 or other SPI derivative Interface, Infiniband, or FC layer.

532. An IP Storage processor having remote direct memory access capability for enabling IP Storage protocols over IP networks, said processor including an IP Storage stack providing IP Storage protocol termination and origination, transporting information in
20 active sessions over IP networks by transporting PDU's specified by the IP storage standard said processor comprising:

- a. An RDMA mechanism for performing RDMA data transfer;
- b. At least one IP Storage processor engine for processing IP Storage
packets;
- 25 c. An IP Storage session memory for storing session information;
- d. At least one memory controller for controlling memory accesses;
- e. At least one media interface for coupling to at least one network; and

f. A host interface for coupling to at least one host or fabric interface for coupling to a fabric.

533. The IP Storage processor of claim 533 further comprising at least one of:

- a. A packet processor engine for processing packets;
- 5 b. A classification processor for classifying IP packets;
- c. A flow controller for controlling data flow;
- d. A policy processor for applying policies;
- e. A security processor for performing security operations;
- f. A controller for control plane processing;
- 10 g. A packet scheduler for scheduling packets;
- h. A packet memory for storing packets; or
- i. A combination of any of the foregoing;

534. The IP Storage processor of claim 532 wherein said processor operates in multiple stages, including one or more stages of

- 15 a. Receiving incoming IP Storage packets;
- b. Providing security processing for said incoming IP Storage packets if needed;
- c. Classifying said incoming IP Storage packets;
- d. Scheduling IP Storage packets for processing;
- 20 e. Executing data and/or protocol processing operations on IP Storage packets;
- f. Providing direct memory access for transferring data/packets to or from the memory of a system external to said processor;

g. Executing protocol processing operations on data or commands forming IP Storage packets;

h. Providing security processing for outgoing IP Storage packets if needed;

5 i. Transmitting outgoing IP Storage packets onto a network; or

j. A combination of any of the foregoing; and

Each of said stages is capable of operating on different IP Storage packets concurrently.

10 535. A TCP/IP processor having transport level RDMA capability for enabling TCP over IP networks, said processor including a TCP/IP stack providing TCP/IP protocol termination and origination, said processor comprising:

a. An RDMA mechanism for performing RDMA data transfer;

b. At least one TCP/IP processor engine for processing IP packets;

c. A session memory for storing session information;

d. At least one memory controller for controlling memory accesses;

15 e. At least one media interface for coupling to at least one network; and

f. A host interface for coupling to at least one host or fabric interface for coupling to a fabric.

536. The TCP/IP processor of claim 535 further comprising at least one of:

a. A packet processor engine for processing packets;

20 b. A classification processor for classifying IP packets;

c. A flow controller for controlling data flow;

d. A policy processor for applying policies;

e. A security processor for performing security operations;

- f. A controller for control plane processing;
- g. A packet scheduler for scheduling packets;
- h. A packet memory for storing packets;
- i. A connection manager or session controller for managing TCP/IP
- 5 sessions; or.
- j. A combination of any of the foregoing;

537. A TCP/IP processor having transport level RDMA capability for enabling TCP over IP networks, said processor including a TCP/IP stack providing TCP/IP protocol termination and origination, said stack providing an interface to sockets layer functions in a

10 host processor to transport data traffic, said processor comprising:

- a. An RDMA mechanism for performing RDMA data transfer;
- b. At least one TCP/IP processor engine for processing IP packets;
- c. A session memory for storing session information;
- d. At least one memory controller for controlling memory accesses;
- 15 e. At least one media interface for coupling to at least one network; and
- f. A host interface for coupling to at least one host or fabric interface for coupling to a fabric.

538. The TCP/IP processor of claim 537 further comprising at least one of:

- a. A packet processor engine for processing packets;
- 20 b. A classification processor for classifying IP packets;
- c. A flow controller for controlling data flow;
- d. A policy processor for applying policies;
- e. A security processor for performing security operations;

171

- f. A controller for control plane processing;
- g. A packet scheduler for scheduling packets;
- h. A packet memory for storing packets;
- i. A connection manager or session controller for managing TCP/IP
- 5 sessions; or
- j. A combination of any of the foregoing;

539. The TCP/IP processor of claim 537 wherein said processor operates in multiple stages, including one or more stages of

- a. Receiving incoming IP packets;
- 10 b. Providing security processing for said incoming IP packets if necessary;
- c. Classifying said incoming IP packets;
- d. Scheduling IP packets for processing;
- e. Executing data and/or protocol processing operations on IP packets;
- 15 f. Providing direct memory access for transferring data/packets to or from the memory of a system external to said processor;
- g. Executing protocol processing operations on data or commands forming IP packets;
- h. Providing security processing for outgoing IP packets if necessary;
- 20 i. Transmitting outgoing IP packets onto a network; or
- j. A combination of any of the foregoing; and
- i. Each of said stages is capable of operating on different IP packets concurrently.

540. The TCP/IP processor of claim 535 wherein said processor operates in multiple stages, including one or more stages of

- a. Receiving incoming IP packets;
 - b. Providing security processing for said incoming IP packets if
5 necessary;
 - c. Classifying said incoming IP packets;
 - d. Scheduling IP packets for processing;
 - e. Executing data and/or protocol processing operations on IP packets;
 - f. Providing direct memory access for transferring data/packets to or
10 from the memory of a system external to said processor;
 - g. Executing protocol processing operations on data or commands forming IP packets;
 - h. Providing security processing for outgoing IP packets if necessary;
 - i. Transmitting outgoing IP packets onto a network; or
 - 15 j. A combination of any of the foregoing; and
- i. Each of said stages is capable of operating on different IP packets concurrently.

541. The switching system of claim 35 wherein said processor operates on said packets to apply an access control, intrusion detection, bandwidth monitoring, bandwidth
20 management, traffic shaping, security, virus detection, anti-spam, quality of service, encryption, decryption, LUN masking, zoning, multi-pathing, link aggregation or virtualization function or policy or a combination of any of the foregoing.

542. A network comprising one or more system, wherein said one or more system is a server, a host bus adapter, a switch, a switch line card, a gateway, a line card of a
25 gateway, a storage area network appliance, a line card of an appliance, a storage system or a line card of a storage system or a combination of any of the foregoing, said one or more

system comprising a hardware processor having transport layer RDMA capability for enabling data transfer using TCP or other session oriented protocols over IP networks, said processor being programmable and comprising a deep packet classification and/or policy processing engine, used by the said system to enable end to end network management for storage and/or non-storage data networks, said processor applying policies on a per packet, per flow, per command basis, or a combination of per packet, or per flow, or per command basis.

543. An IP processor, capable of transport level RDMA, for enabling TCP or other session oriented protocols over IP networks, said processor comprising:

- 10 a. at least one packet processor for processing IP packets;
- b. a session memory for storing IP session information;
- c. at least one memory controller for controlling memory accesses;
- d. at least one media interface for coupling to at least one network; and
- e. a host interface for coupling to at least one host or a fabric interface
- 15 for coupling to a fabric;
- i. wherein said processor operates in multiple stages, including one or more of the stages of (1) receiving incoming IP packets; (2) providing security for processing said incoming IP packets if needed; (3) classifying said incoming IP packets; (4) scheduling IP packets for processing; (5) executing data and/or processing operations on
- 20 IP Packets; (6) providing direct memory access for transferring data/packets to or from the memory of a system external to said processor; (7) executing protocol processing operations on data or commands forming IP packets; (8) providing processing security for outgoing IP packets, if needed; or (9) transmitting IP packets onto a network; or any combination of the foregoing; and
- 25 ii. each of said stages is capable of operating on different IP packets concurrently.

544. The IP processor of claim 543 wherein each stage of said IP processor may take a different length of time to perform its function than one or more of the other stages of said IP processor.

5 545. A TCP/IP processor for enabling TCP over IP networks, said processor including a TCP/IP stack providing TCP/IP protocol termination and origination, said processor comprising:

- a. At least one TCP/IP processor engine for processing IP packets;
- b. A session memory for storing session information;
- c. At least one memory controller for controlling memory accesses;
- 10 d. At least one media interface for coupling to at least one network; and
- e. A host interface for coupling to at least one host or fabric interface for coupling to a fabric.

546. The TCP/IP processor of claim 545 further comprising at least one of:

- a. A packet processor engine for processing packets;
- 15 b. A classification processor for classifying IP packets;
- c. A flow controller for controlling data flow;
- d. A policy processor for applying policies;
- e. A security processor for performing security operations;
- f. A controller for control plane processing;
- 20 g. A packet scheduler for scheduling packets;
- h. A packet memory for storing packets;
- i. A connection manager or session controller for managing TCP/IP sessions; or

175

- j. A combination of any of the foregoing.

547. A TCP/IP processor for enabling TCP over IP networks, said processor including a TCP/IP stack providing TCP/IP protocol termination and origination, said stack providing an interface to sockets layer functions in a host processor to transport data traffic, said processor comprising:

- 5
- a. At least one TCP/IP processor engine for processing IP packets;
 - b. A session memory for storing session information;
 - c. At least one memory controller for controlling memory accesses;
 - d. At least one media interface for coupling to at least one network; and
 - 10 e. A host interface for coupling to at least one host or fabric interface for coupling to a fabric.

548. The TCP/IP processor of claim 547 further comprising at least one of:

- 15
- a. A packet processor engine for processing packets;
 - b. A classification processor for classifying IP packets;
 - c. A flow controller for controlling data flow;
 - d. A policy processor for applying policies;
 - e. A security processor for performing security operations;
 - f. A controller for control plane processing;
 - g. A packet scheduler for scheduling packets;
 - 20 h. A packet memory for storing packets;
 - i. A connection manager or session controller for managing TCP/IP sessions; or
 - j. A combination of any of the foregoing.

549. An IP Storage processor for enabling IP Storage protocols over IP networks, said processor including an IP Storage stack providing IP Storage protocol termination and origination, transporting information in active sessions over IP networks by transporting PDU's specified by the IP storage standard said processor comprising:

- 5 a. At least one IP Storage processor engine for processing IP Storage packets;
- b. An IP Storage session memory for storing session information;
- c. At least one memory controller for controlling memory accesses;
- d. At least one media interface for coupling to at least one network; and
- 10 e. A host interface for coupling to at least one host or fabric interface for coupling to a fabric.

550. The IP Storage processor of claim 549 further comprising at least one of:

- a. A packet processor engine for processing packets;
- b. A classification processor for classifying IP packets;
- 15 c. A flow controller for controlling data flow;
- d. A policy processor for applying policies;
- e. A security processor for performing security operations;
- f. A controller for control plane processing;
- g. A packet scheduler for scheduling packets;
- 20 h. A packet memory for storing packets; or.
- i. A combination of any of the foregoing.

551. The TCP/IP processor of claim 544 wherein said processor operates in multiple stages, including one or more stages of

177

- a. Receiving incoming IP packets;
- b. Providing security processing for said incoming IP packets if
5 necessary;
- c. Classifying said incoming IP packets;
- d. Scheduling IP packets for processing;
- e. Executing data and/or protocol processing operations on IP packets;
- f. Providing direct memory access for transferring data/packets to or
from the memory of a system external to said processor;
- g. Executing protocol processing operations on data or commands
10 forming IP packets;
- h. Providing security processing for outgoing IP packets if necessary;
- i. Transmitting outgoing IP packets onto a network; or
- j. A combination of any of the foregoing; and
- i. Each of said stages is capable of operating on different IP
15 packets concurrently.

552. The TCP/IP processor of claim 547 wherein said processor operates in multiple stages, including one or more stages of

- a. Receiving incoming IP packets;
- b. Providing security processing for said incoming IP packets if
20 necessary;
- c. Classifying said incoming IP packets;
- d. Scheduling IP packets for processing;
- e. Executing data and/or protocol processing operations on IP packets;

f. Providing direct memory access for transferring data/packets to or from the memory of a system external to said processor;

g. Executing protocol processing operations on data or commands forming IP packets;

5 h. Providing security processing for outgoing IP packets if necessary;

i. Transmitting outgoing IP packets onto a network; or

j. A combination of any of the foregoing; and

i. Each of said stages is capable of operating on different IP packets concurrently.

10 553. The IP Storage processor of claim 549 wherein said processor operates in multiple stages, including one or more stages of

a. Receiving incoming IP Storage packets;

b. Providing security processing for said incoming IP Storage packets if necessary;

15 c. Classifying said incoming IP Storage packets;

d. Scheduling IP packets for processing;

e. Executing data and/or protocol processing operations on IP Storage packets;

20 f. Providing direct memory access for transferring data/packets to or from the memory of a system external to said processor;

g. Executing protocol processing operations on data or commands forming IP packets;

h. Providing security processing for outgoing IP Storage packets if necessary;

- i. Transmitting outgoing IP packets onto a network; or
- j. A combination of any of the foregoing; and

- i. Each of said stages is capable of operating on different IP Storage packets concurrently.

- 5 554. A network comprising one or more system, wherein said one or more system is a server, a host bus adapter, a switch, a switch line card, a gateway, a line card of a gateway, a storage area network appliance, a line card of an appliance, a storage system or a line card of a storage system or a combination of any of the foregoing, said one or more system comprising a hardware processor for enabling data transfer using TCP or other
- 10 session oriented protocols over IP networks, said processor being programmable and comprising a deep packet classification and/or policy processing engine, used by the said system to enable end to end network management for storage and/or non-storage data networks, said processor applying policies on a per packet, per flow, per command basis, or a combination of per packet, or per flow, or per command basis.

Fig. 1 SCSI Architecture Layers

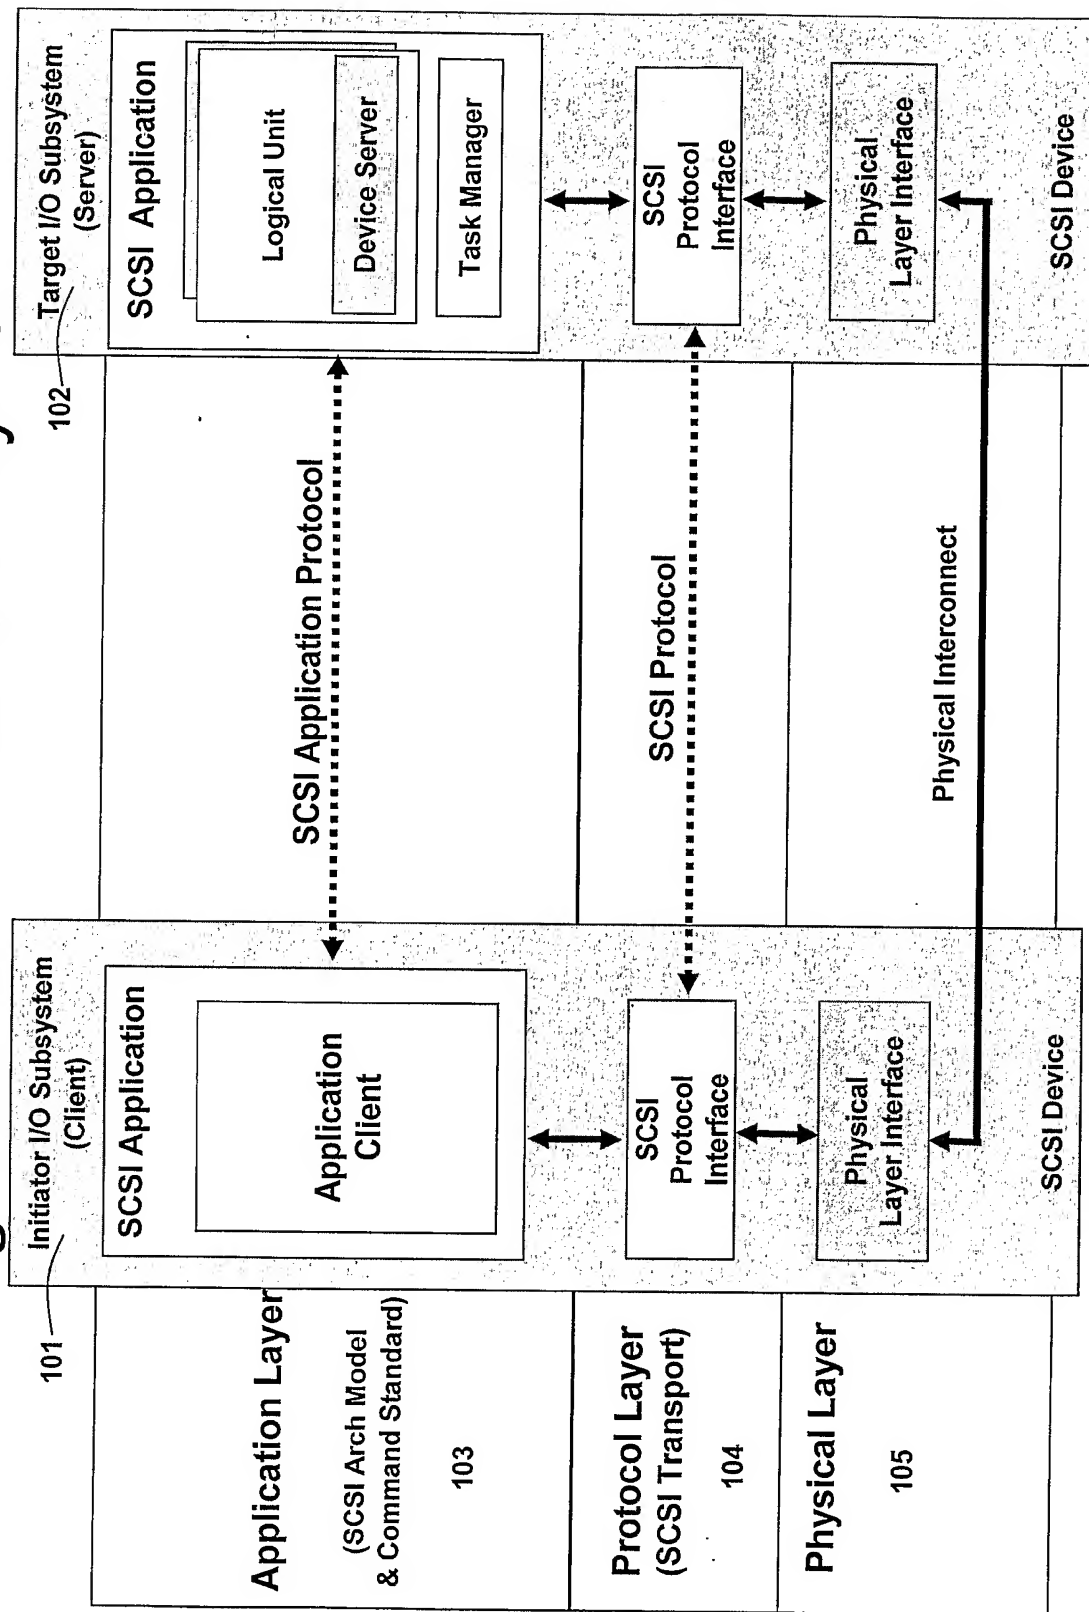


Fig. 2 SCSI Architecture Layers on iSCSI

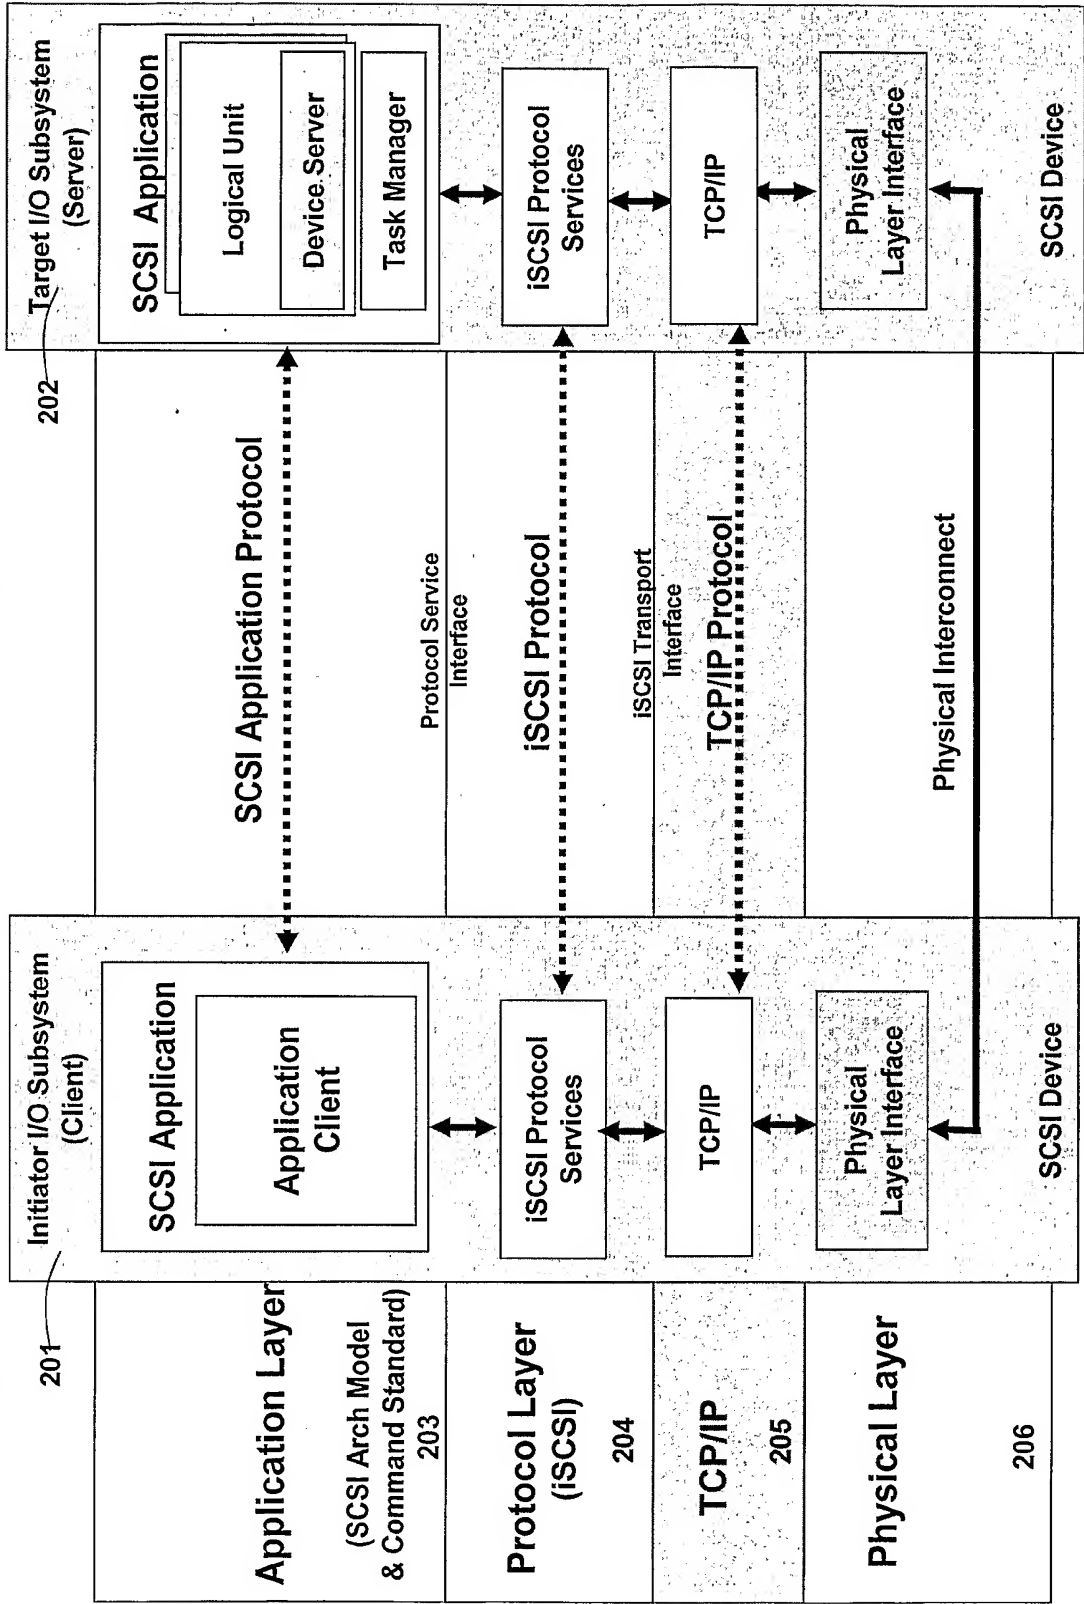


Fig. 3 OSI Stack

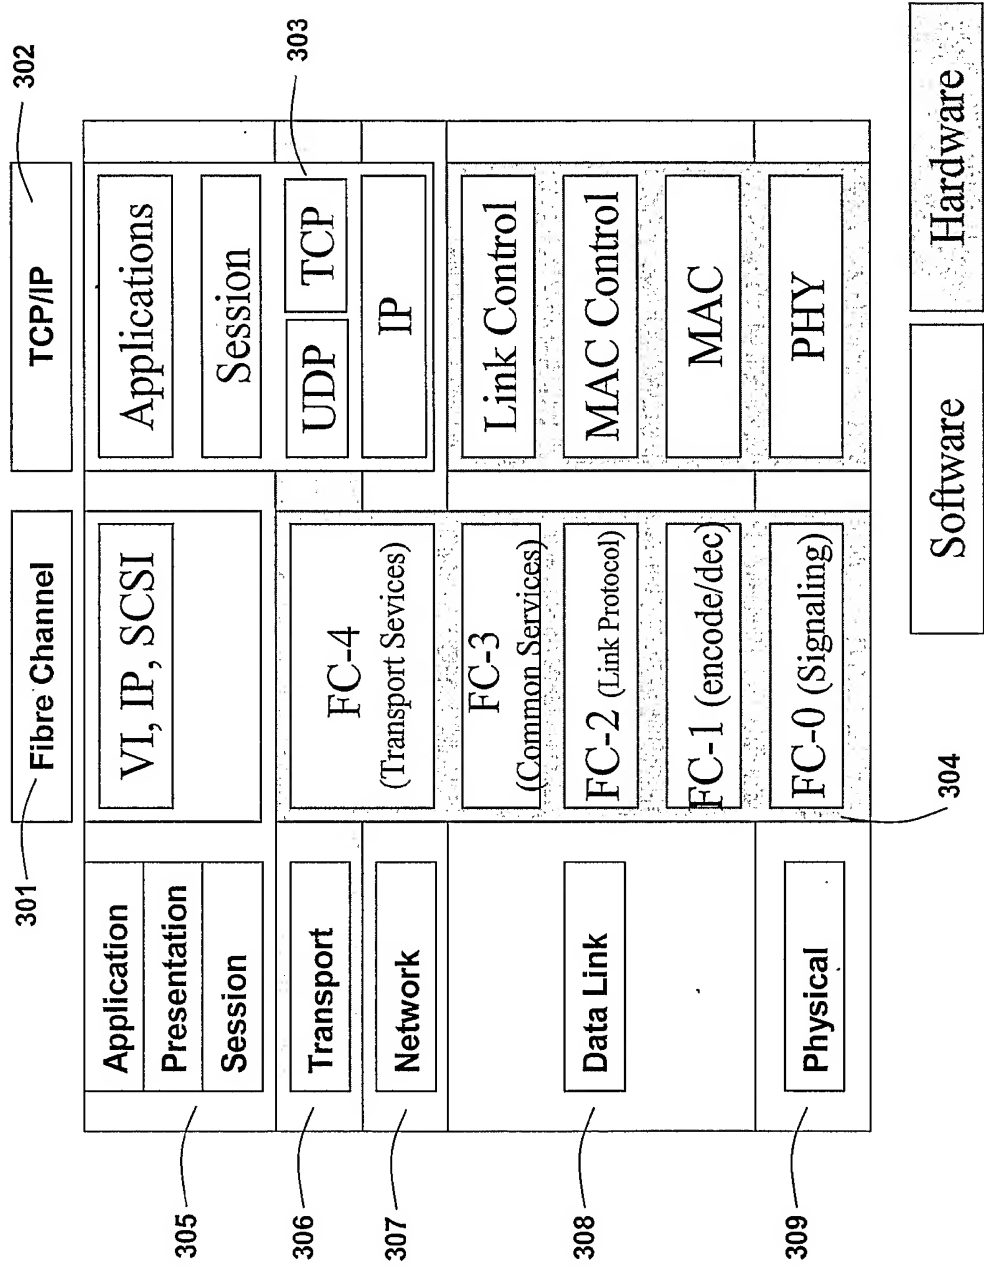


Fig. 4 OSI Stack (w/ HW TCP/IP)

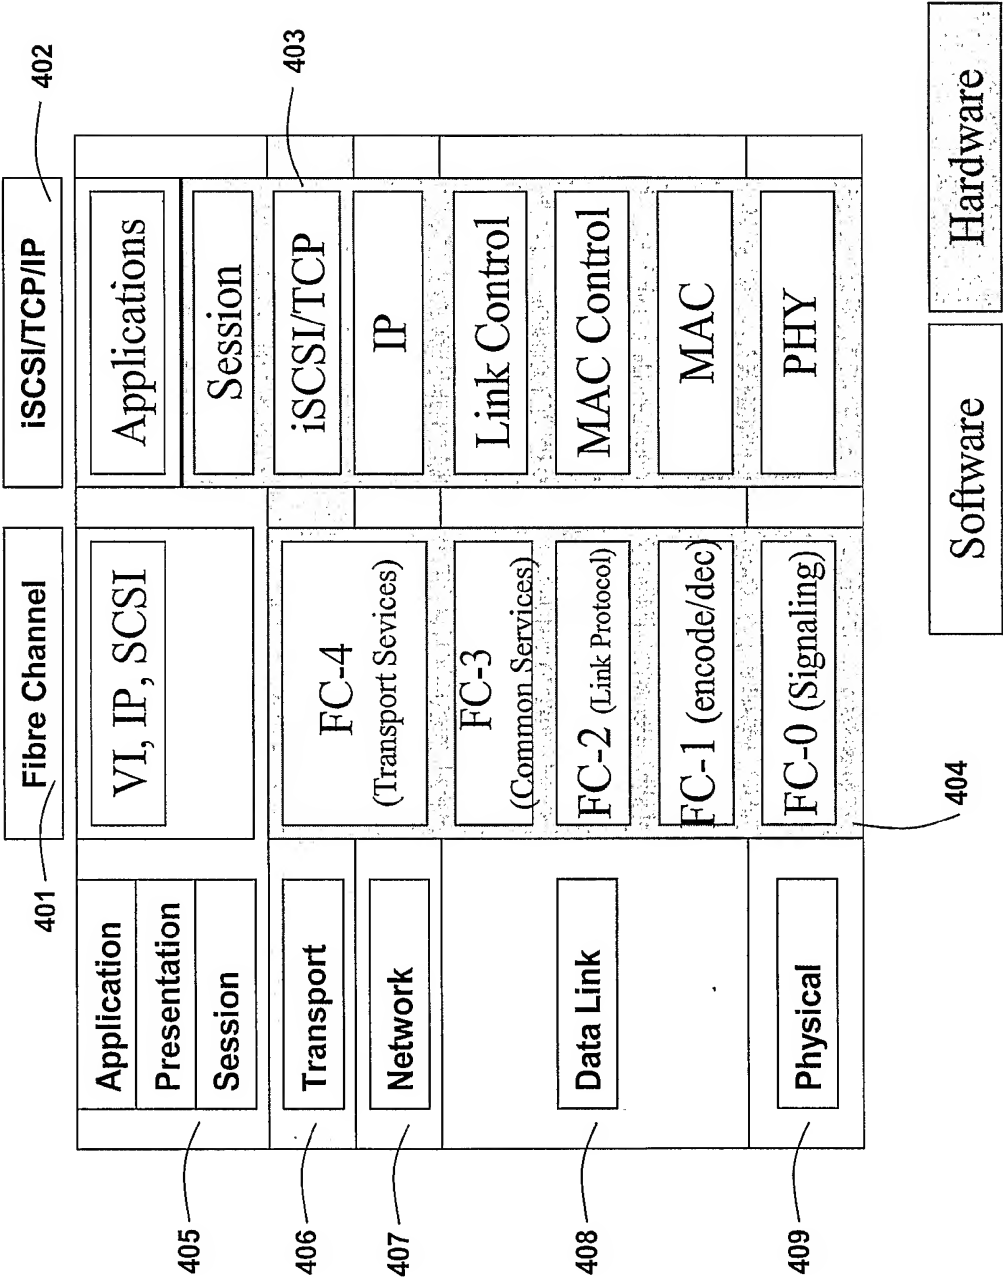


Fig. 5 Host Software Stack (iSCSI)




Fig. 6 S/W TCP Stack Data Transfer




Fig. 7 Remote Direct Memory Access




Fig. 8 Host Software Stack (SCSI)




Fig. 9 iSCSI Stack




Fig. 10 TCP/IP Stack

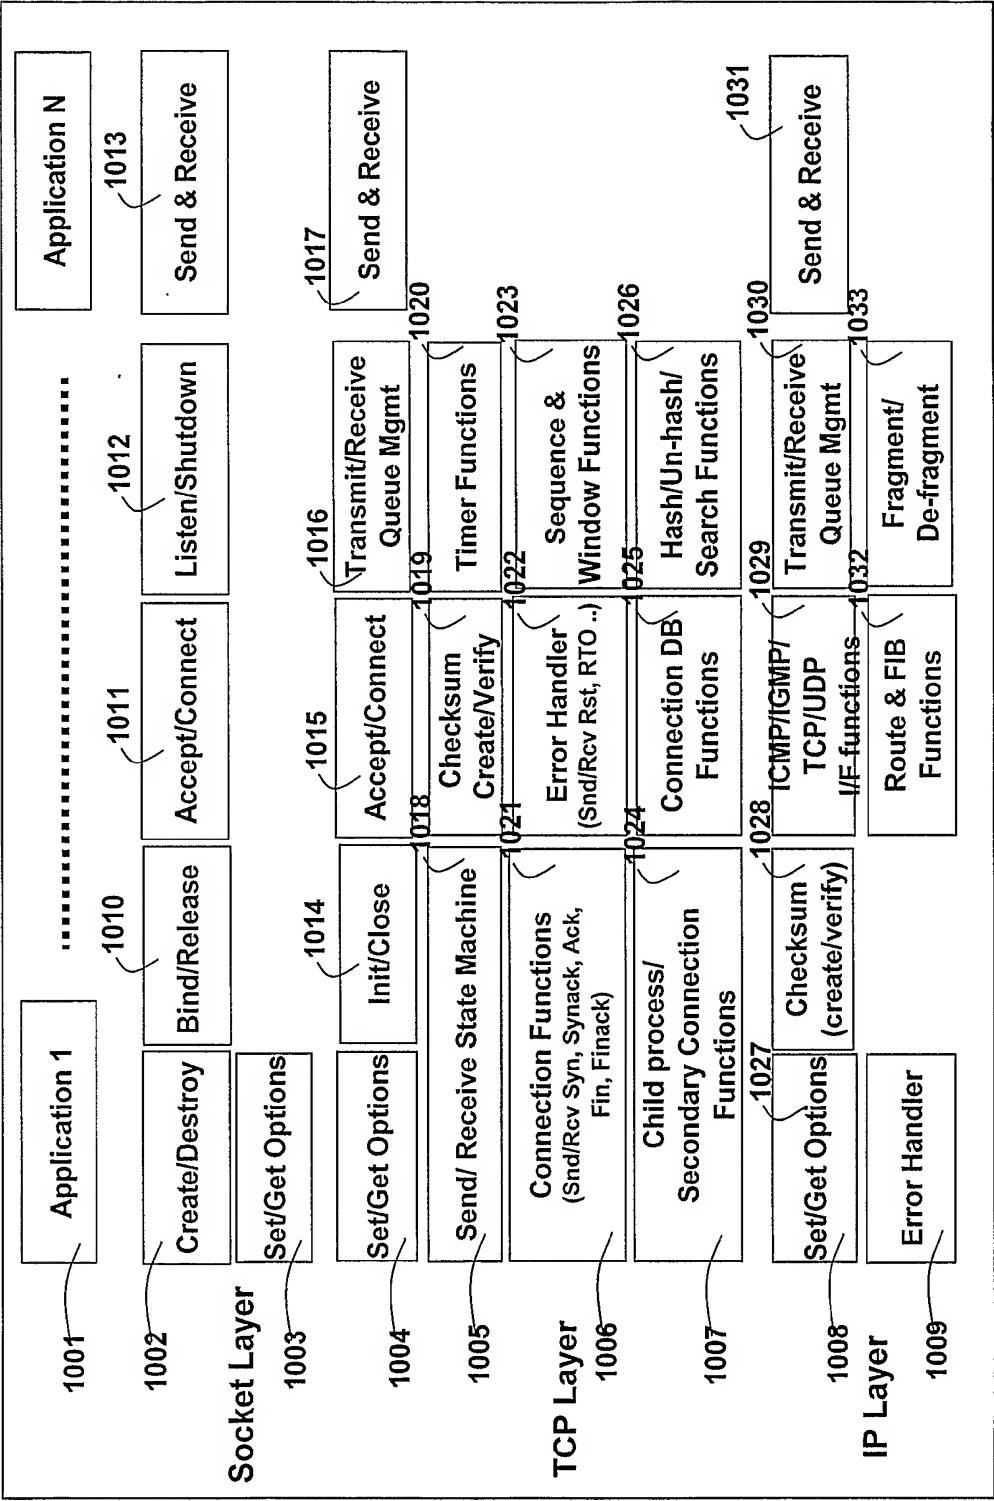


Fig. 11 iSCSI Data Flow




Fig. 12 iSCSI Data Structures




Fig. 13 TCP/IP Session DB Entry

Source_IP	Destination_IP	Source_Port	Destination_Port	Protocol
Connection_ID	TCP_Window_PTR	Window_Size	Sender MSS/MTU	Expected_SQN
TCP_State	TCP_SND_SEQ#	TCP_SEQ#	TCP_ACK#	TCP_RCV_SEQ#
Fragment_PTR	Packet_send_PTR (byte granular)	Packet_Start_Time	Last_PKT_Time	TCP_Parameters (slow start, retransmit..)
Packets_Count	Congestion Window	SSThresh	Receiver MSS	Other Fields
Frame_key	Out of Order PTR	Other Frame Fields		
Hash_Key	NXT_H_Match_PTR	Valid/Invalid		

Fig. 14 iSCSI Session DB Entry

iSCSI_State	Data_SEQ#	Command_SEQ#	Task_TAG	R2T_SEQ#
Status_SEQ#	EXP_Status_SEQ#	EXP_CMD_SEQ#	TARGET_MODE (Solicited or not)	Target_Parameters first burst, max PDU ...
ISID	TSID	Portal Group Tag	Transfer Direction	Other fields
Connection_ID	Command ID	RDMA enabled	RDMA Keys	Other RDMA Param
MDL_List_Pointer	Mem Descriptor	Block size	No. of Blocks (xmitted or rcv'ed)	ULP_Pointers
Hash_Key	NXT_H_Match_PTR	Valid/Invalid		

Fig. 15 iSCSI Session Memory





Fig. 16 IP Network Application Processor






Fig. 17 IP Network Application Processor

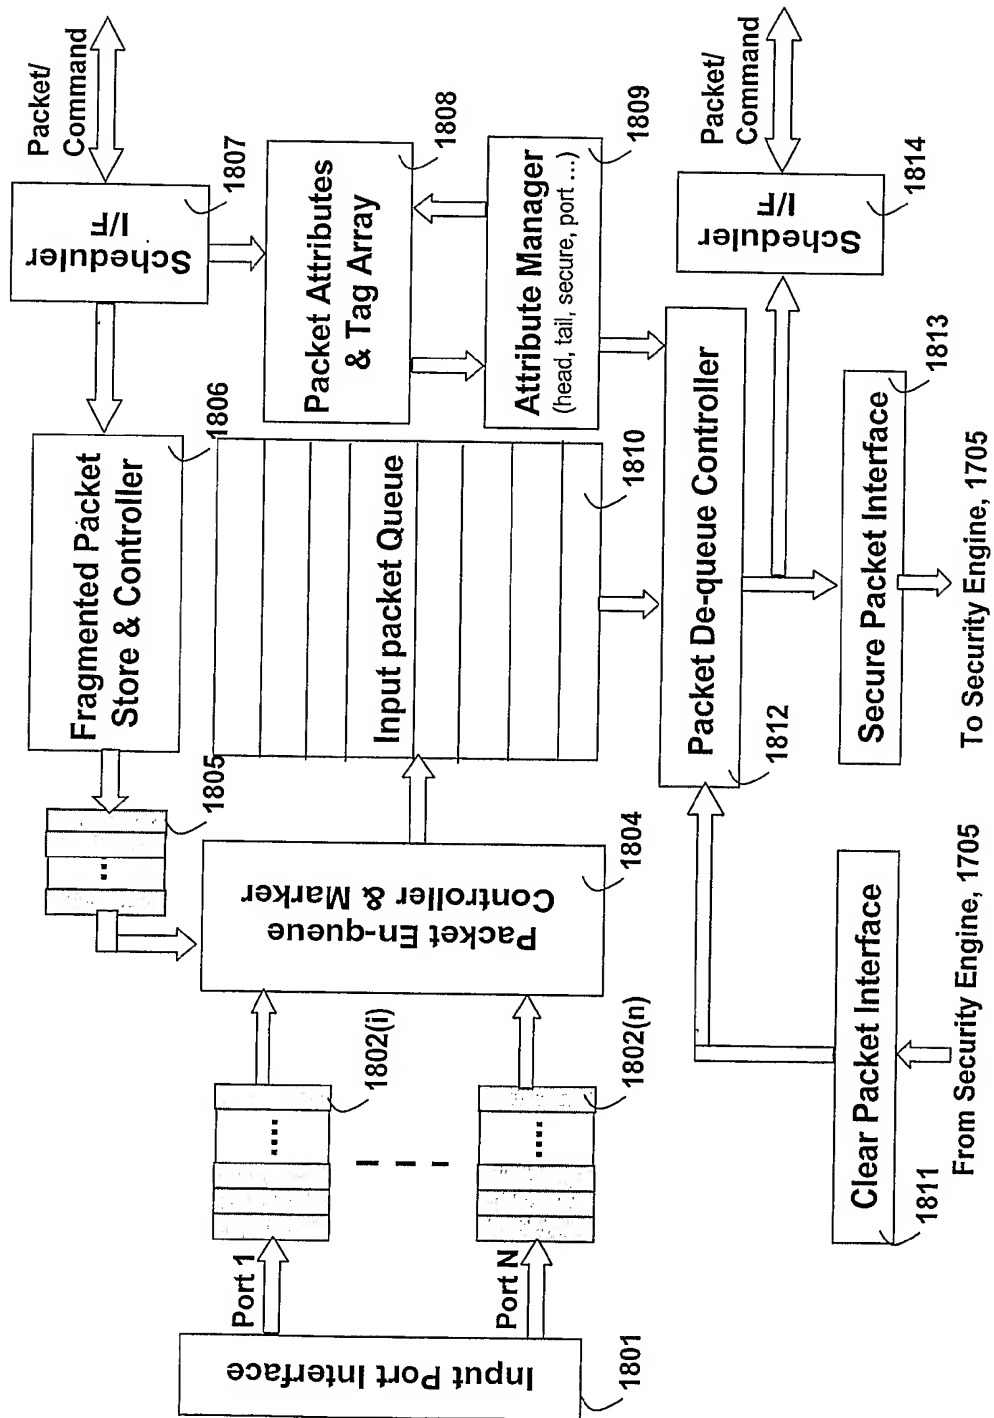
Fig. 18 Input Queue & Controller

Fig. 20 Packet Classification/Policy Engine

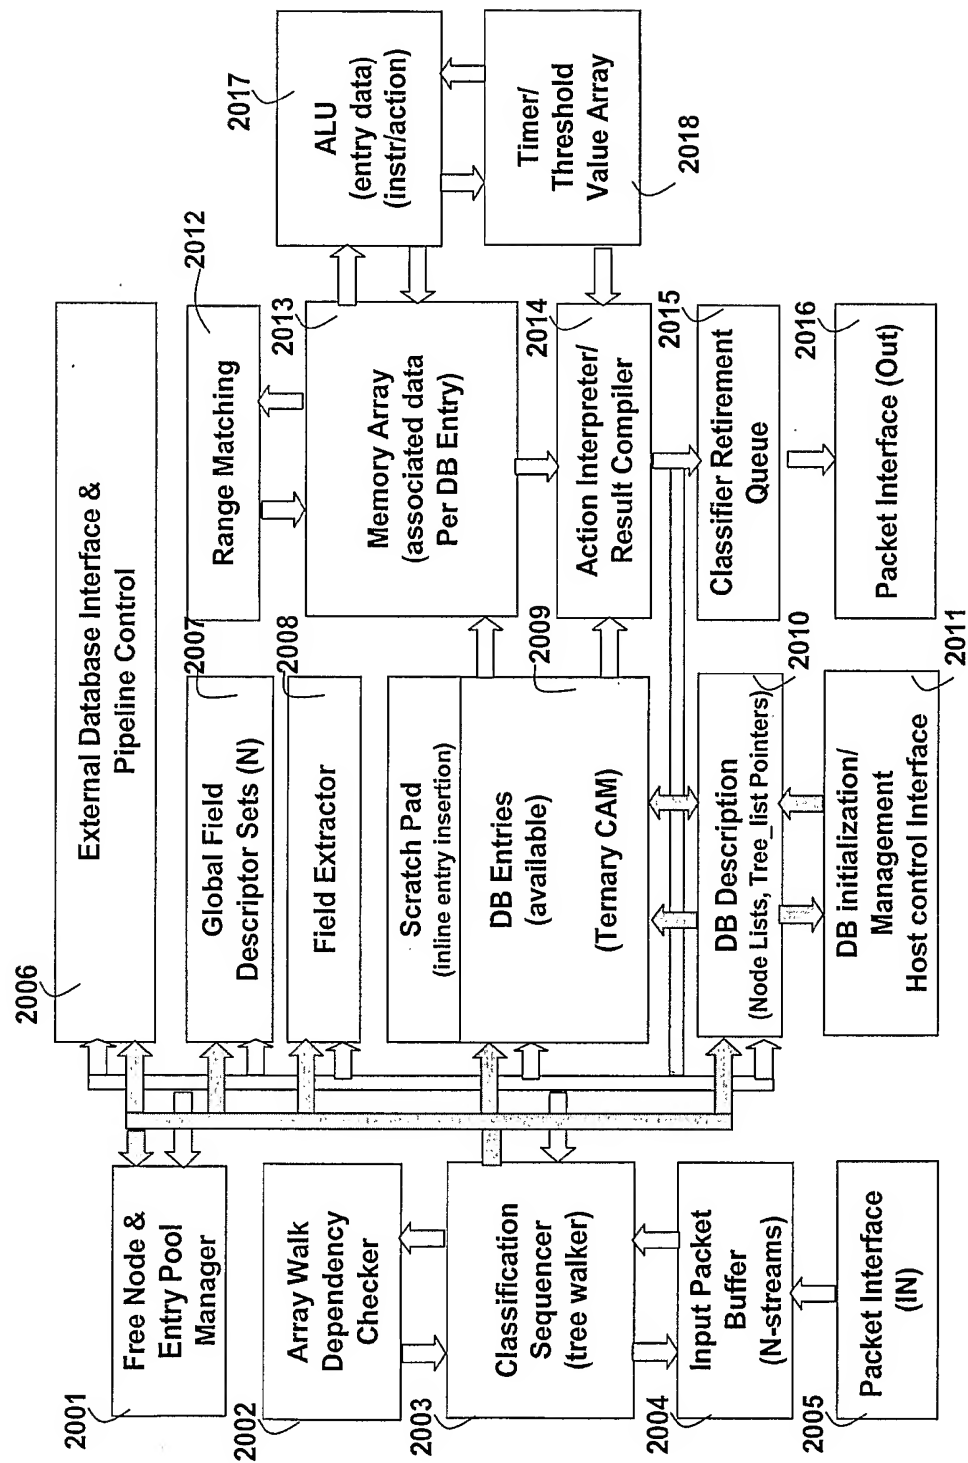


Fig. 21 SAN Packet Processor




Fig. 22 SAN Packet Processor




Fig. 23 TCP/IP Processor




Fig. 24 IP Storage Engine (iSCSI Processor)




Fig. 25 Output Queue

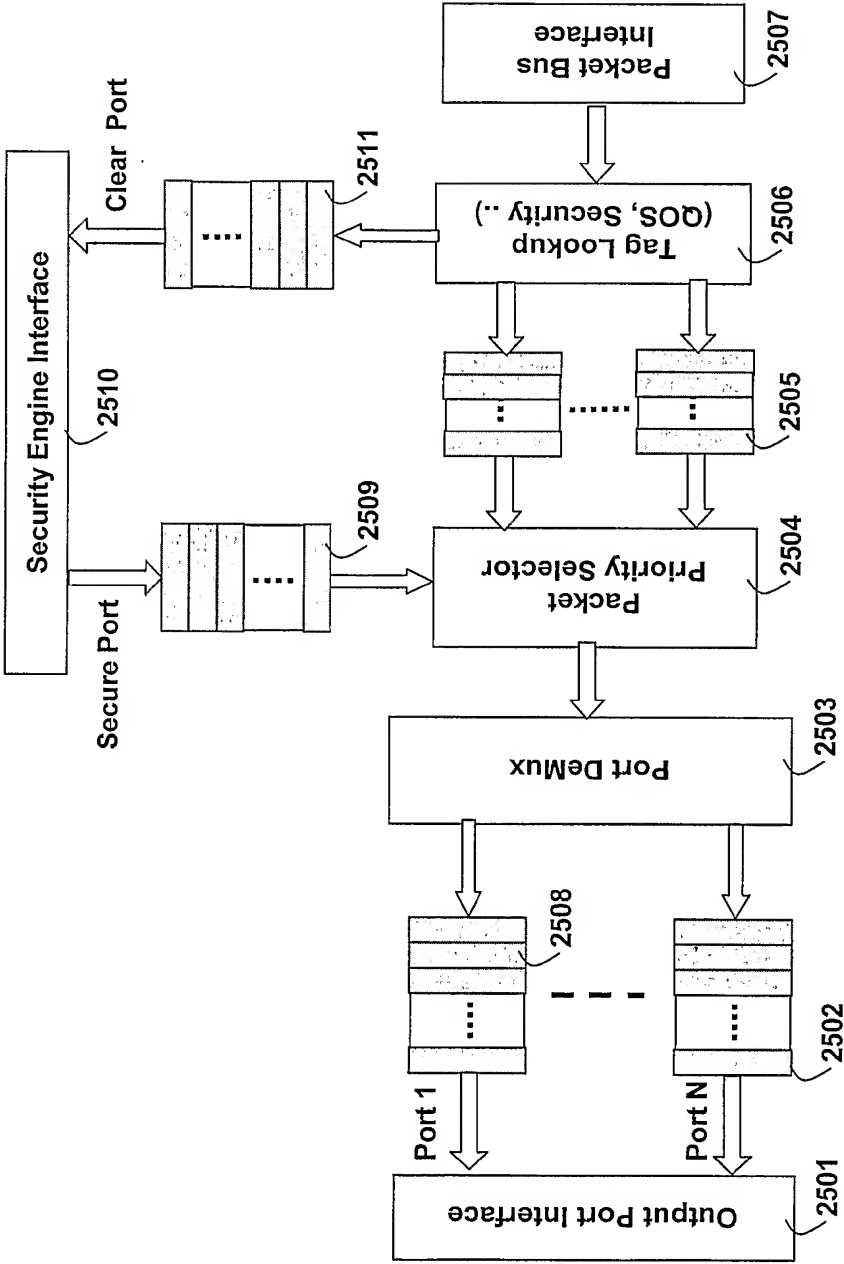


Fig. 26 Storage Flow & RDMA Controller

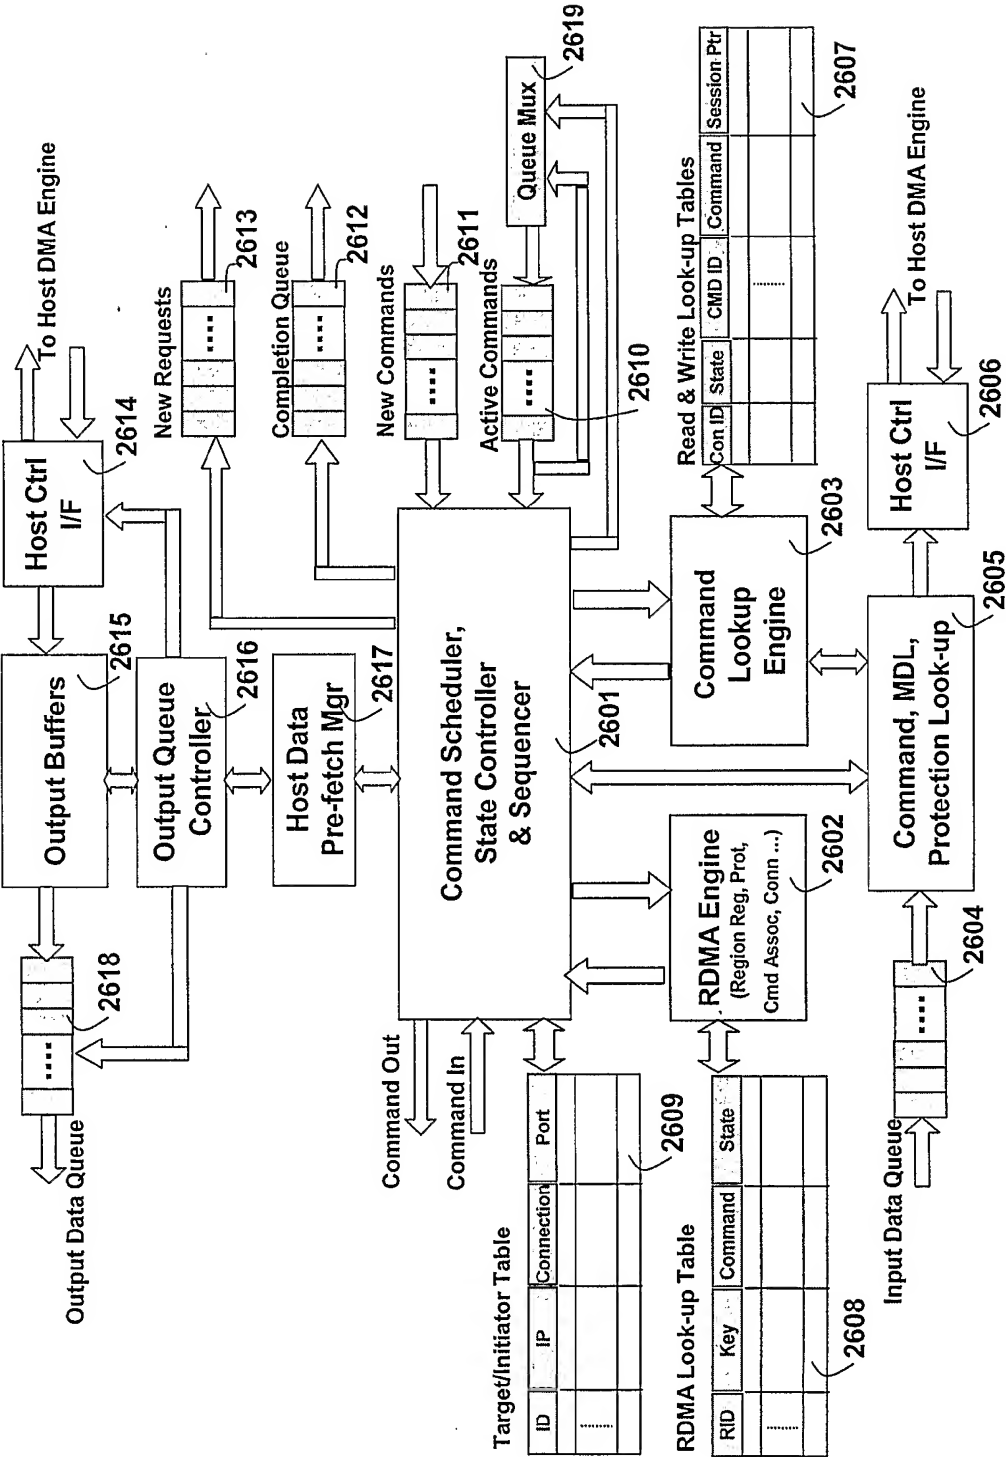


Fig. 27 Host Interface Controller

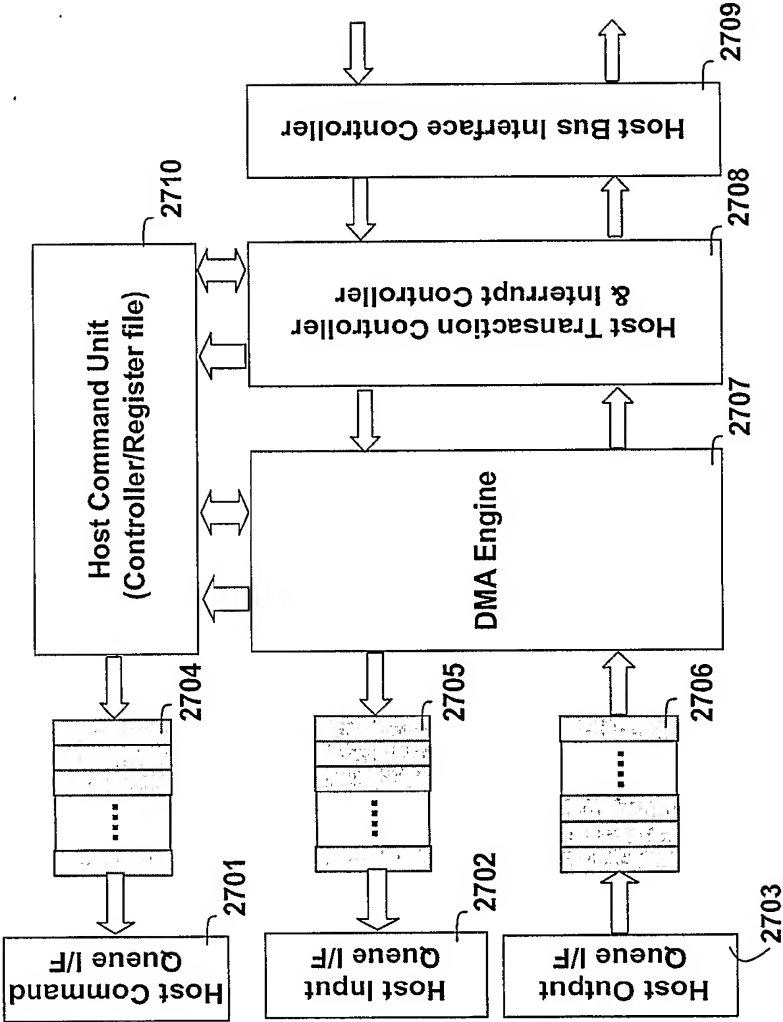


Fig. 28 Security Engine

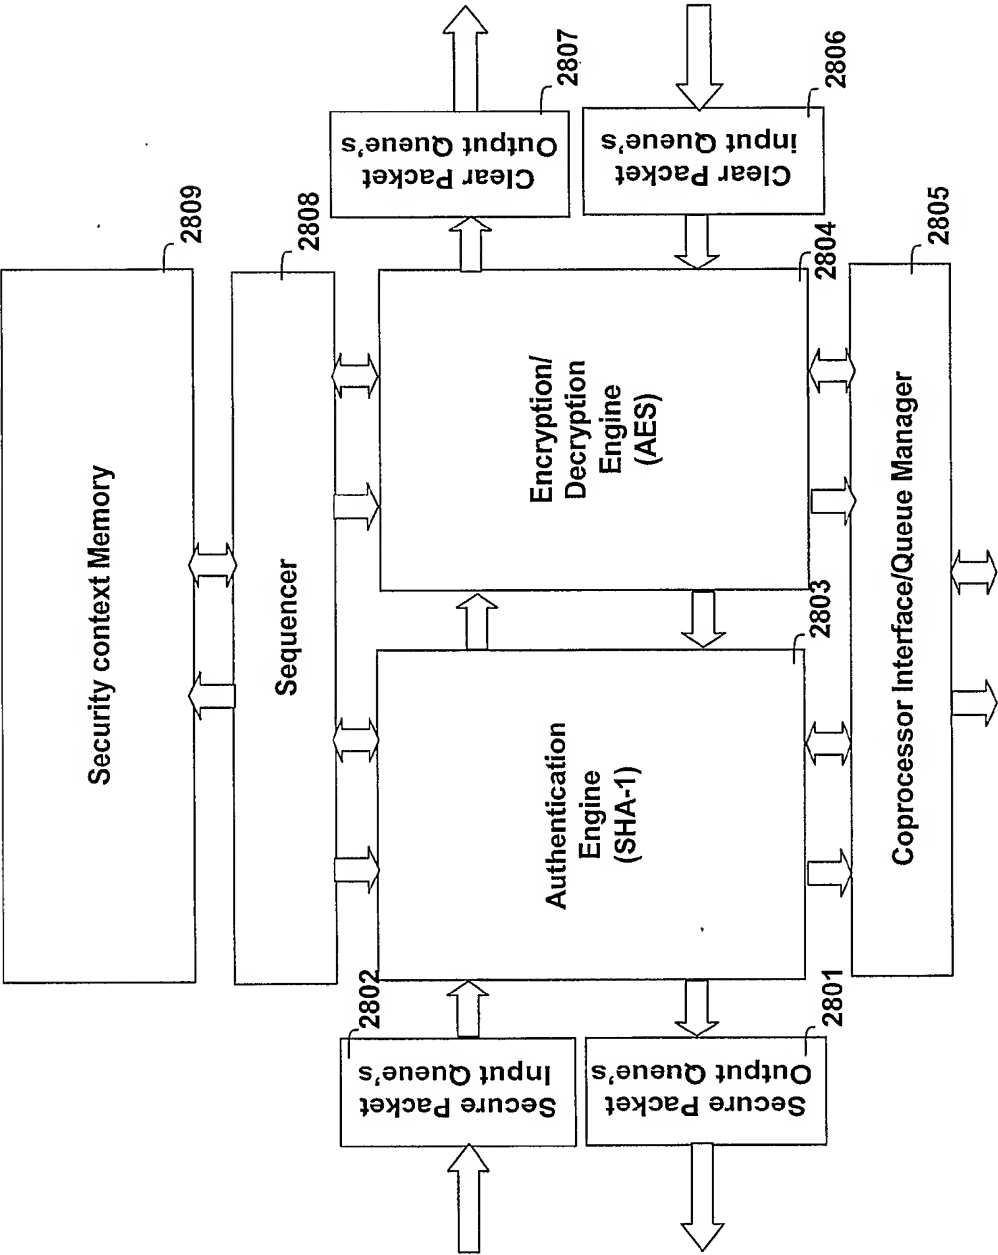


Fig. 29 IP Storage Session Cache & Memory Controller Complex

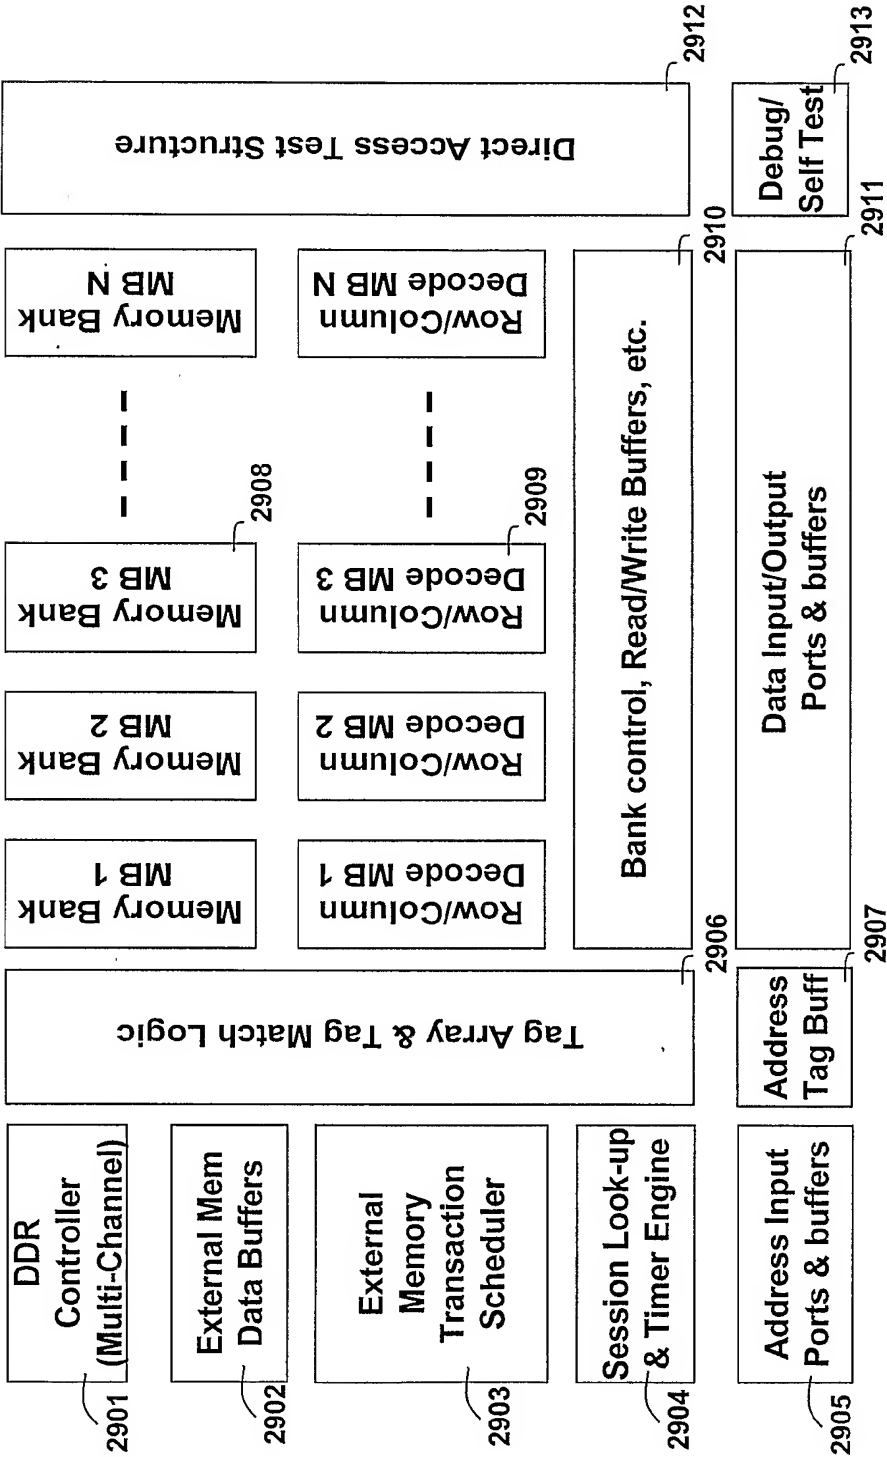


Fig. 30 Classification Engine Data Structure

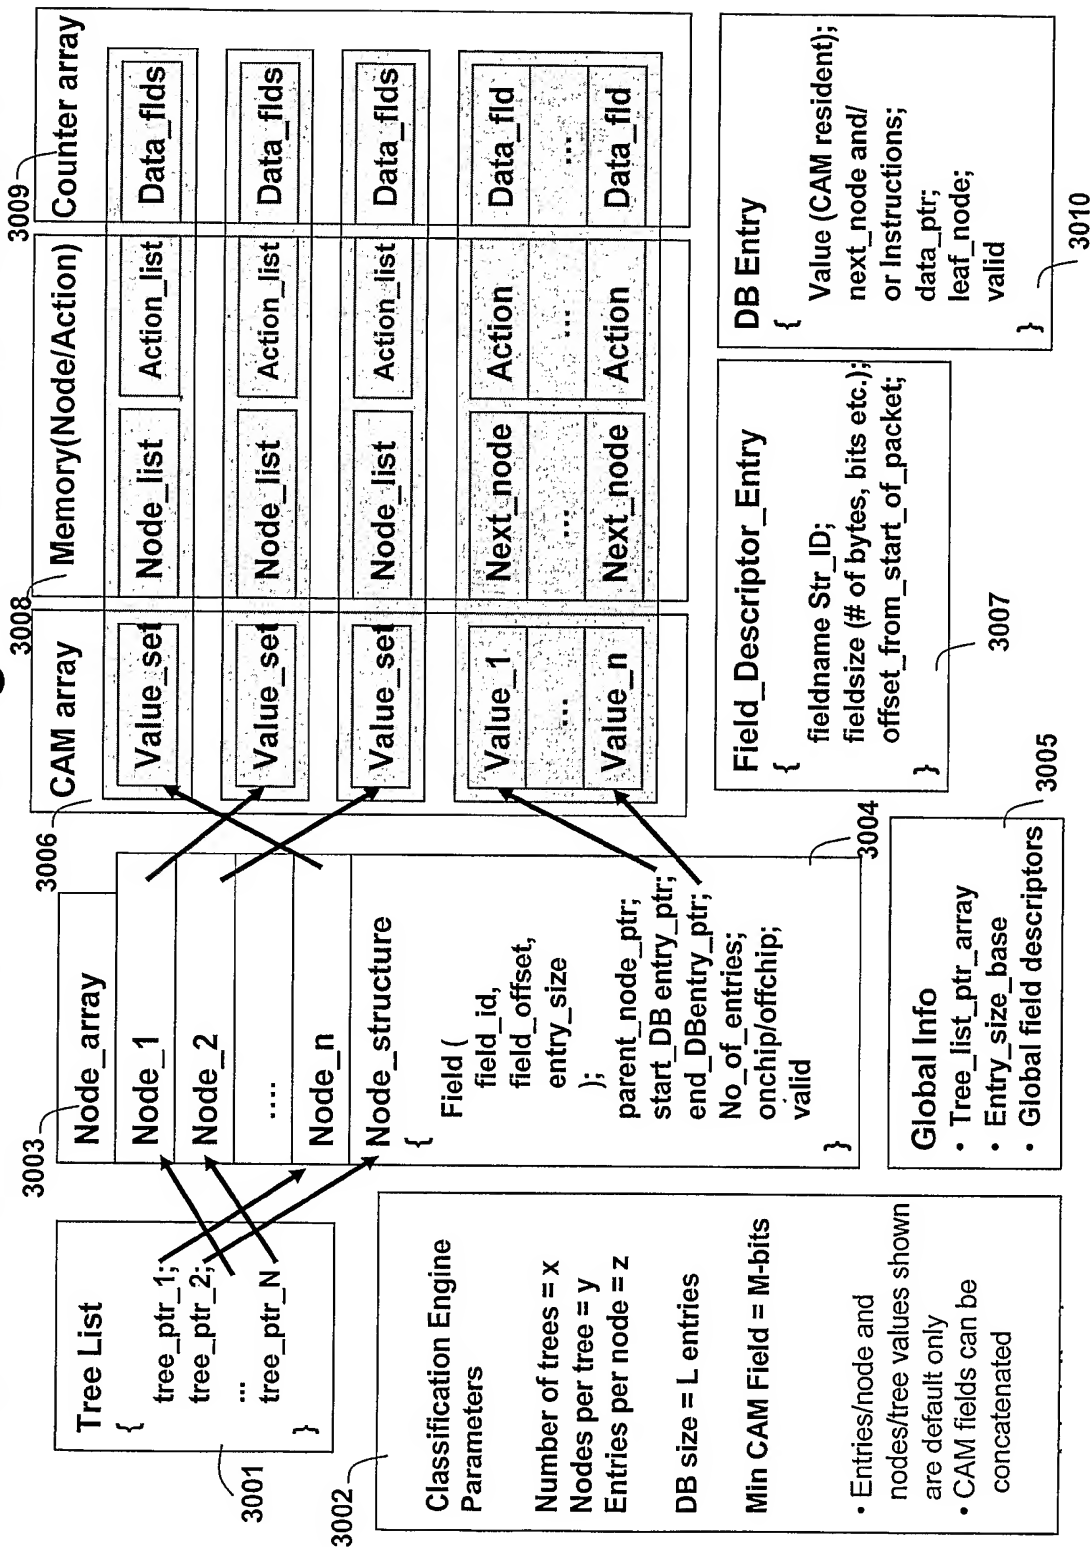


Fig. 31 Read Operation

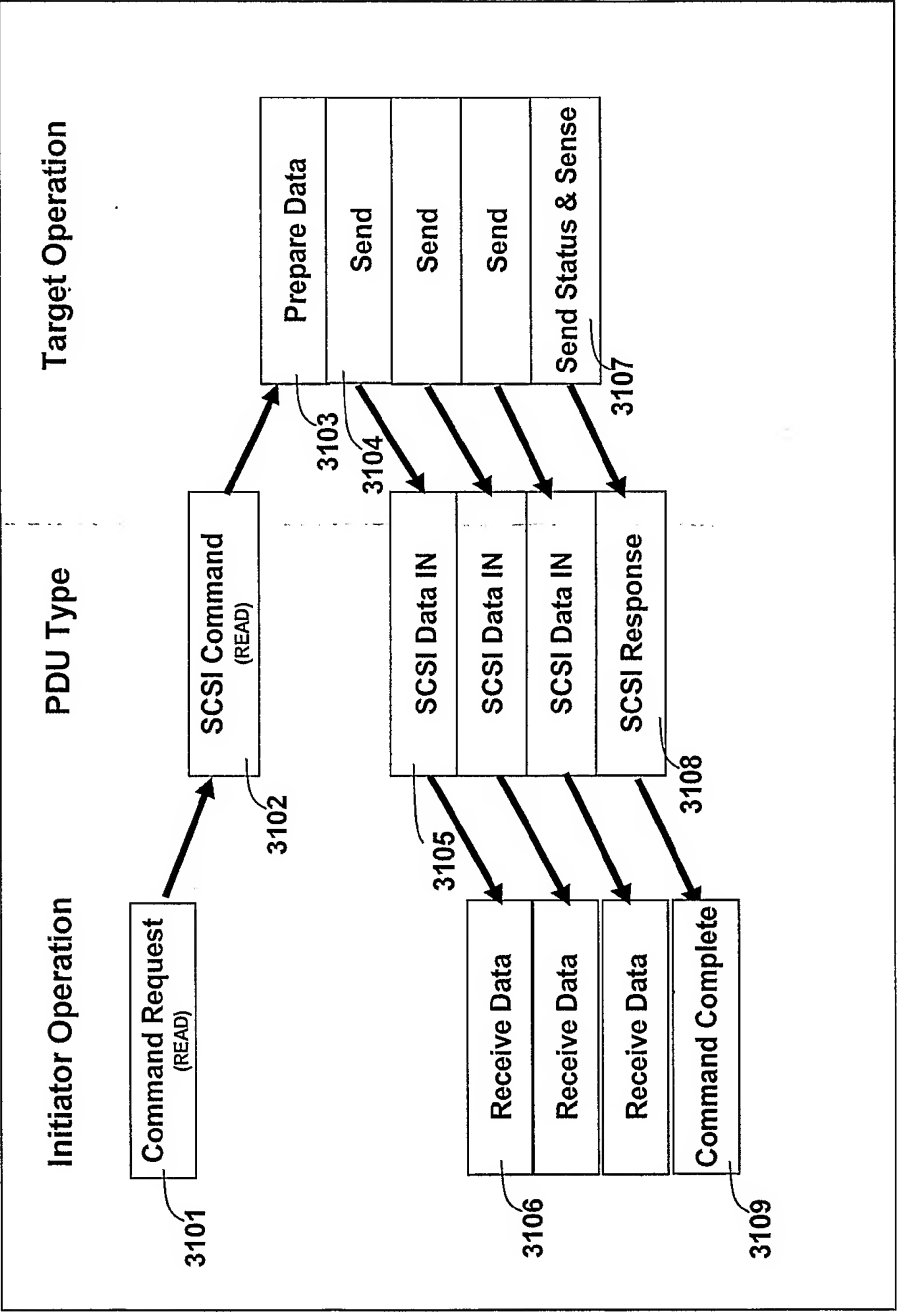


Fig. 32 Read Data Packet Flow

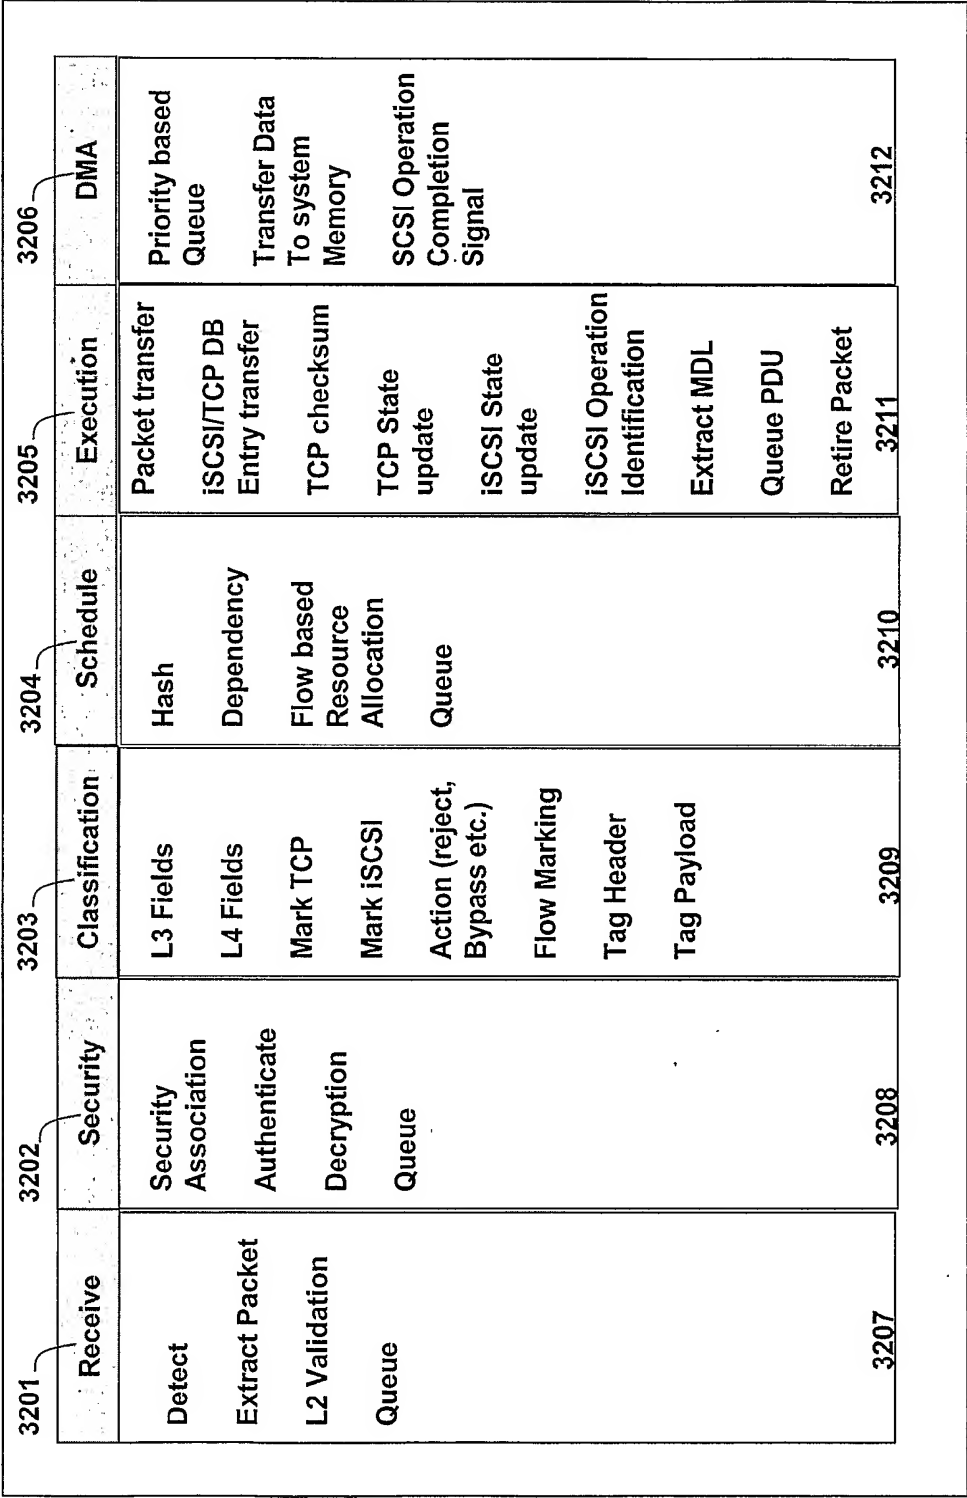


Fig. 33 Write Operation

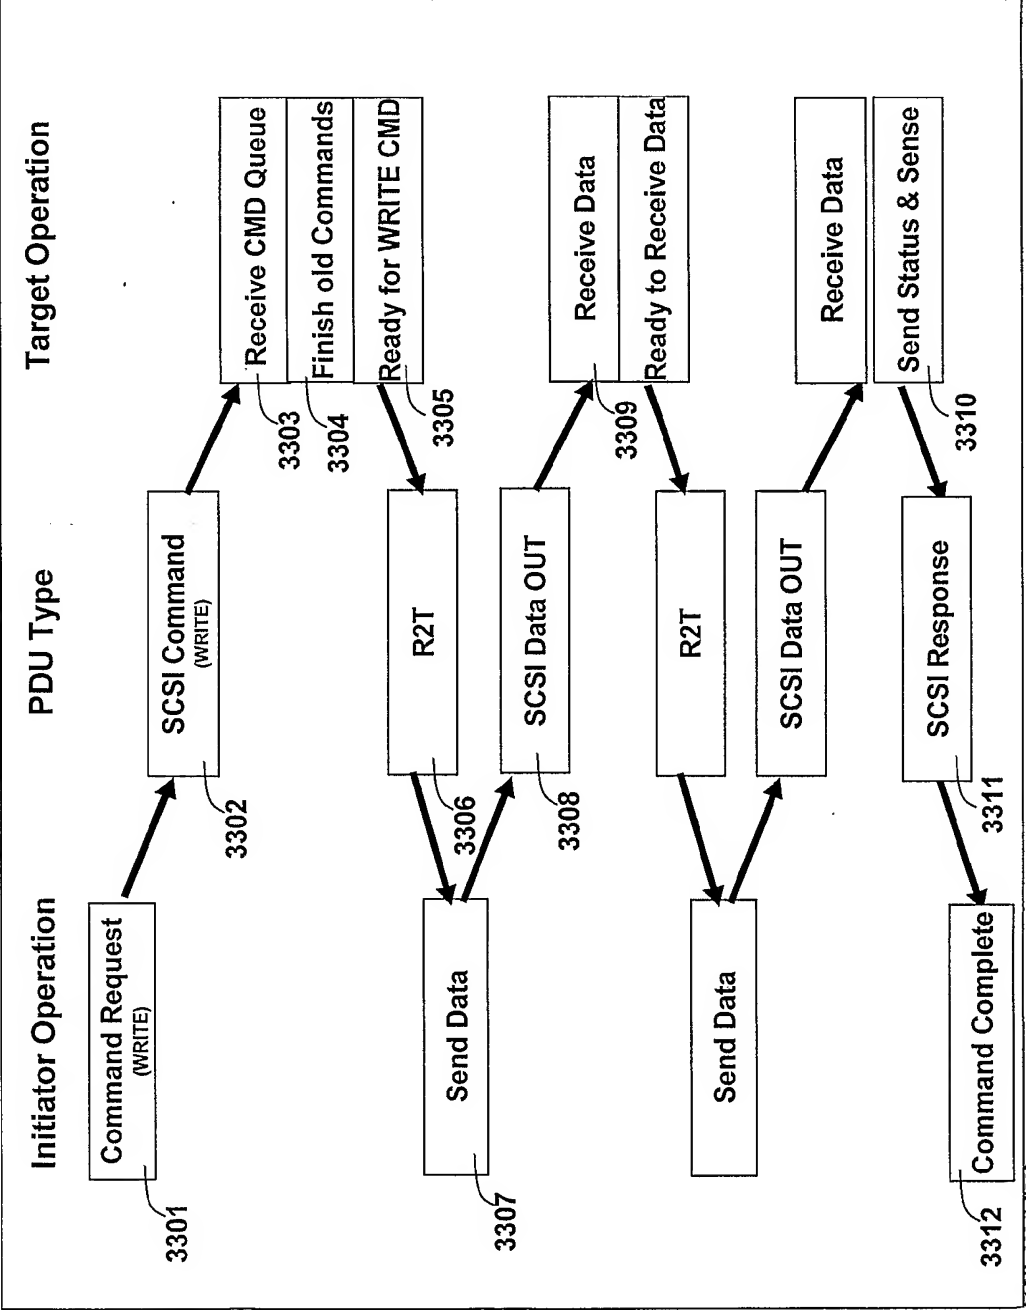


Fig. 34 Write Data Packet Flow

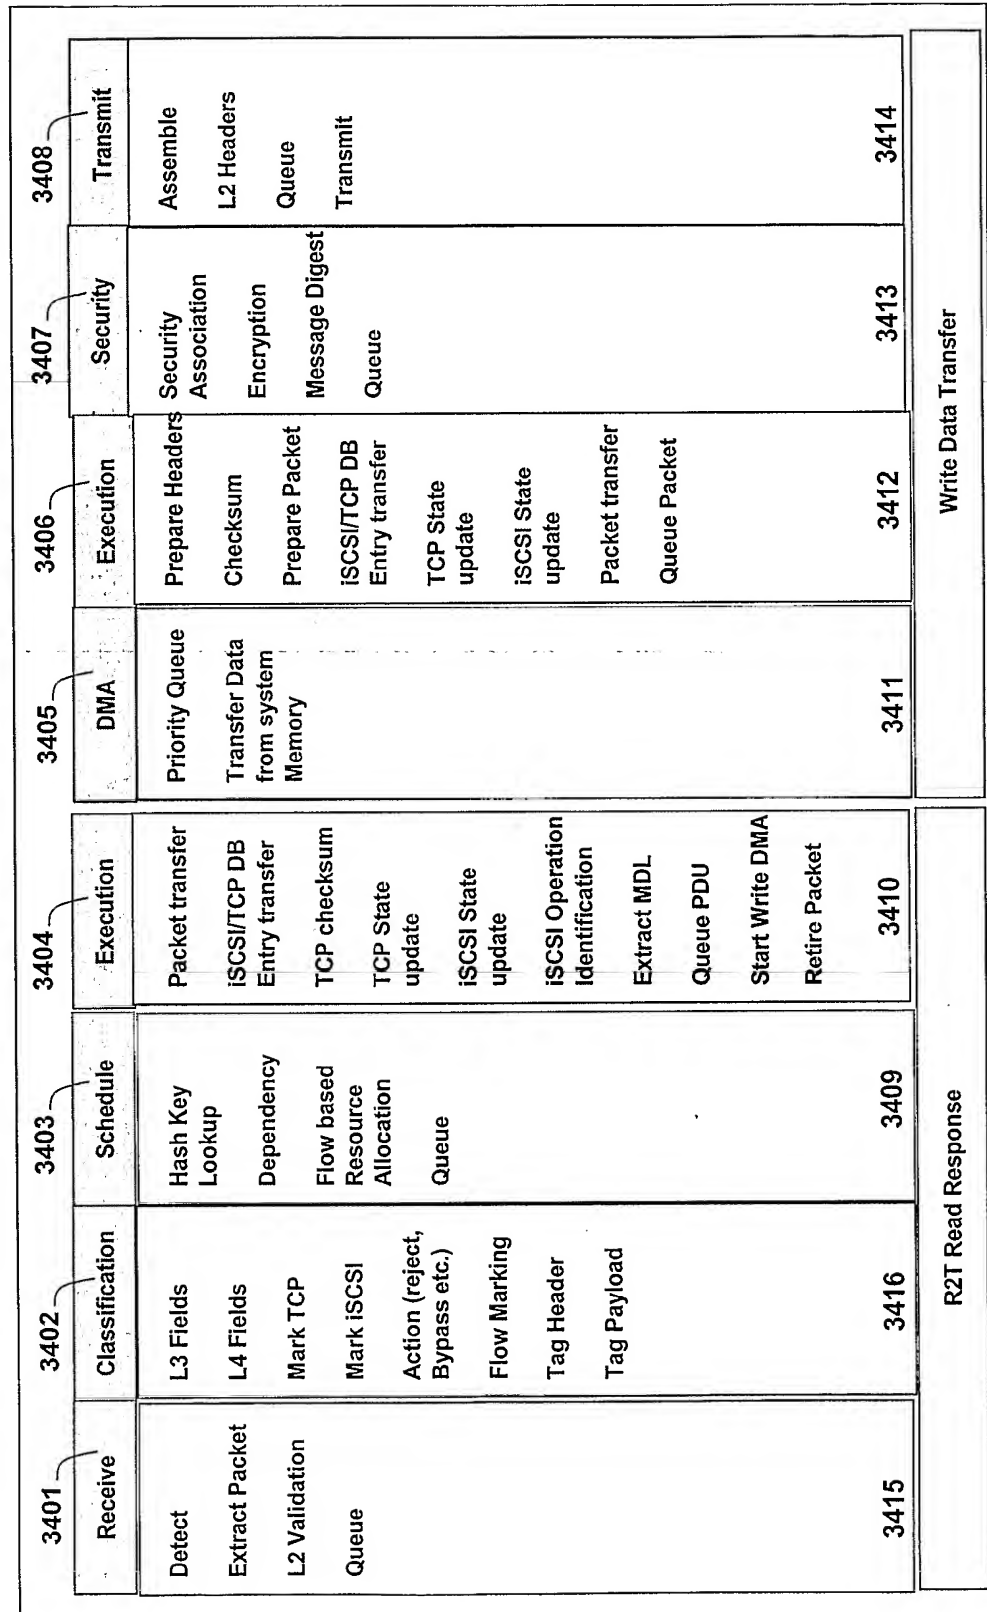


Fig. 35 iSCSI Read using RDMA Write

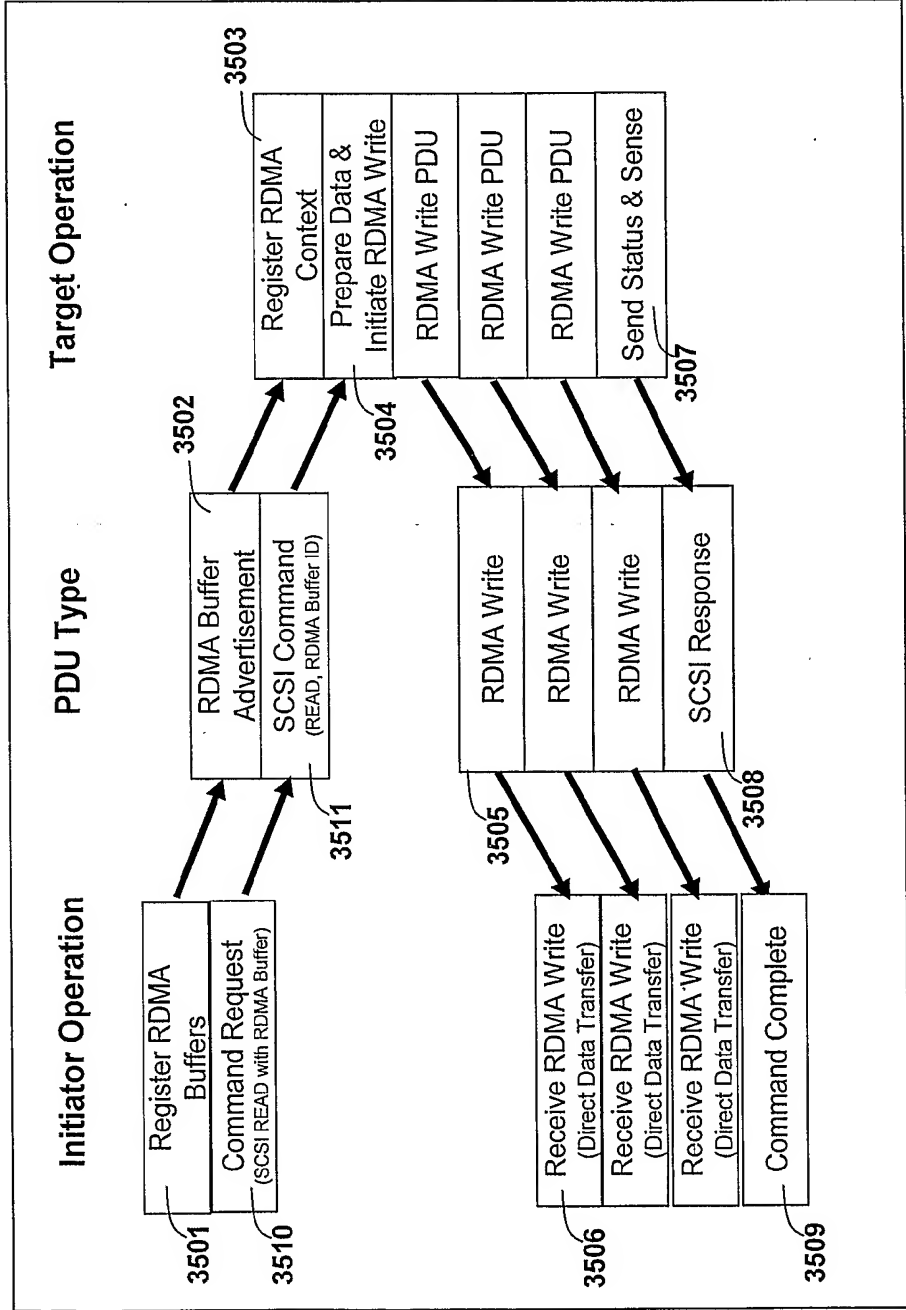


Fig. 36 Read Data Packet Flow (RDMA)

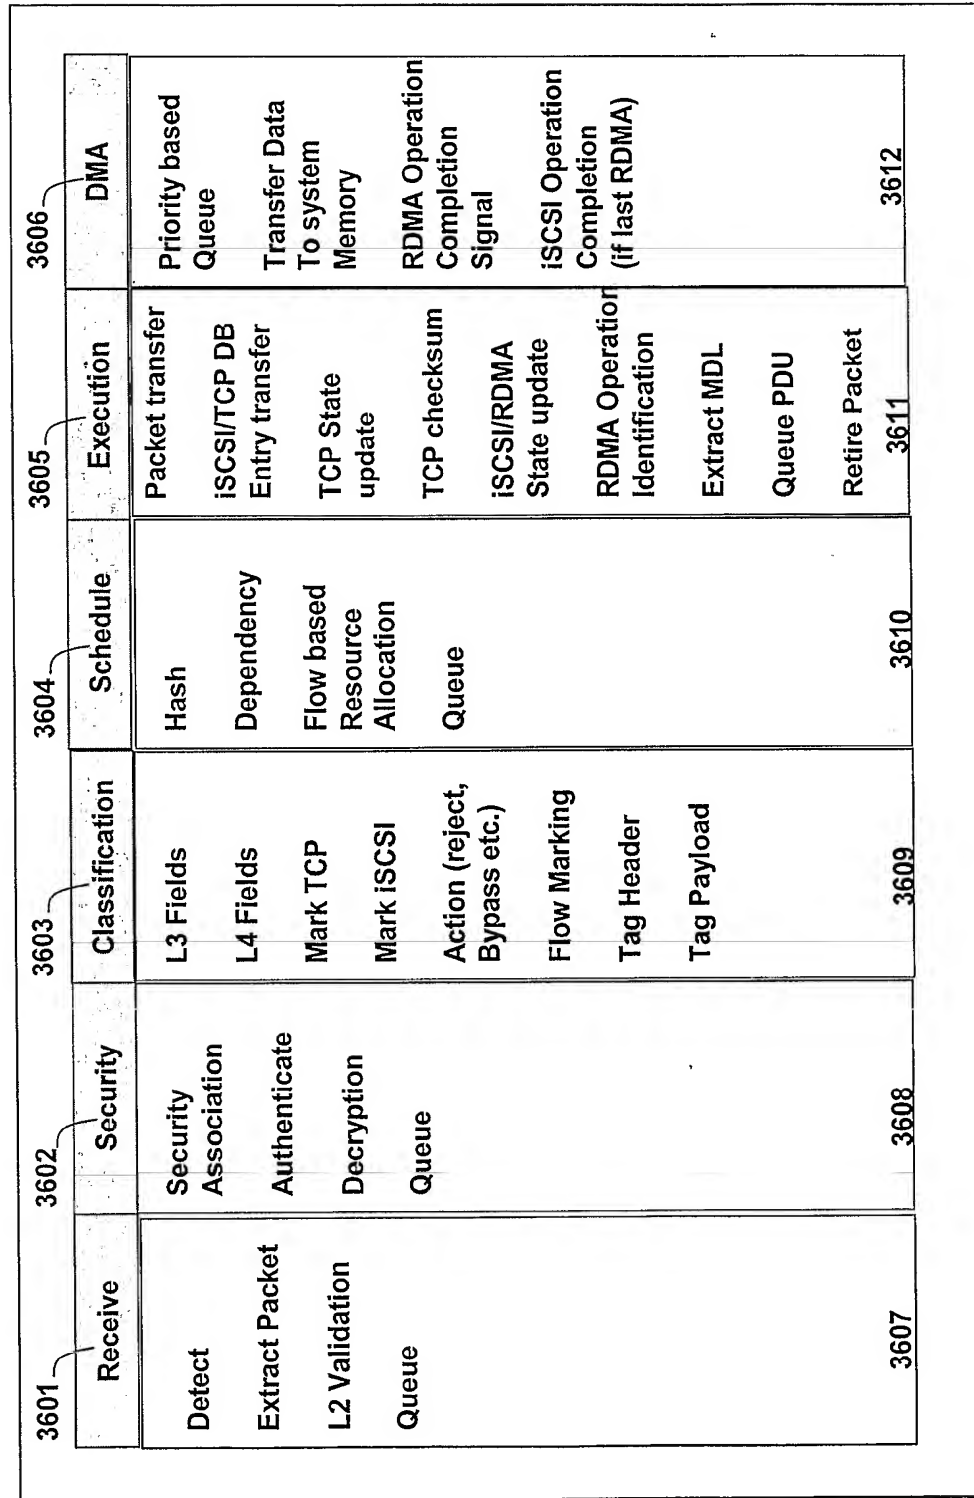


Fig. 37 iSCSI Write using RDMA Read

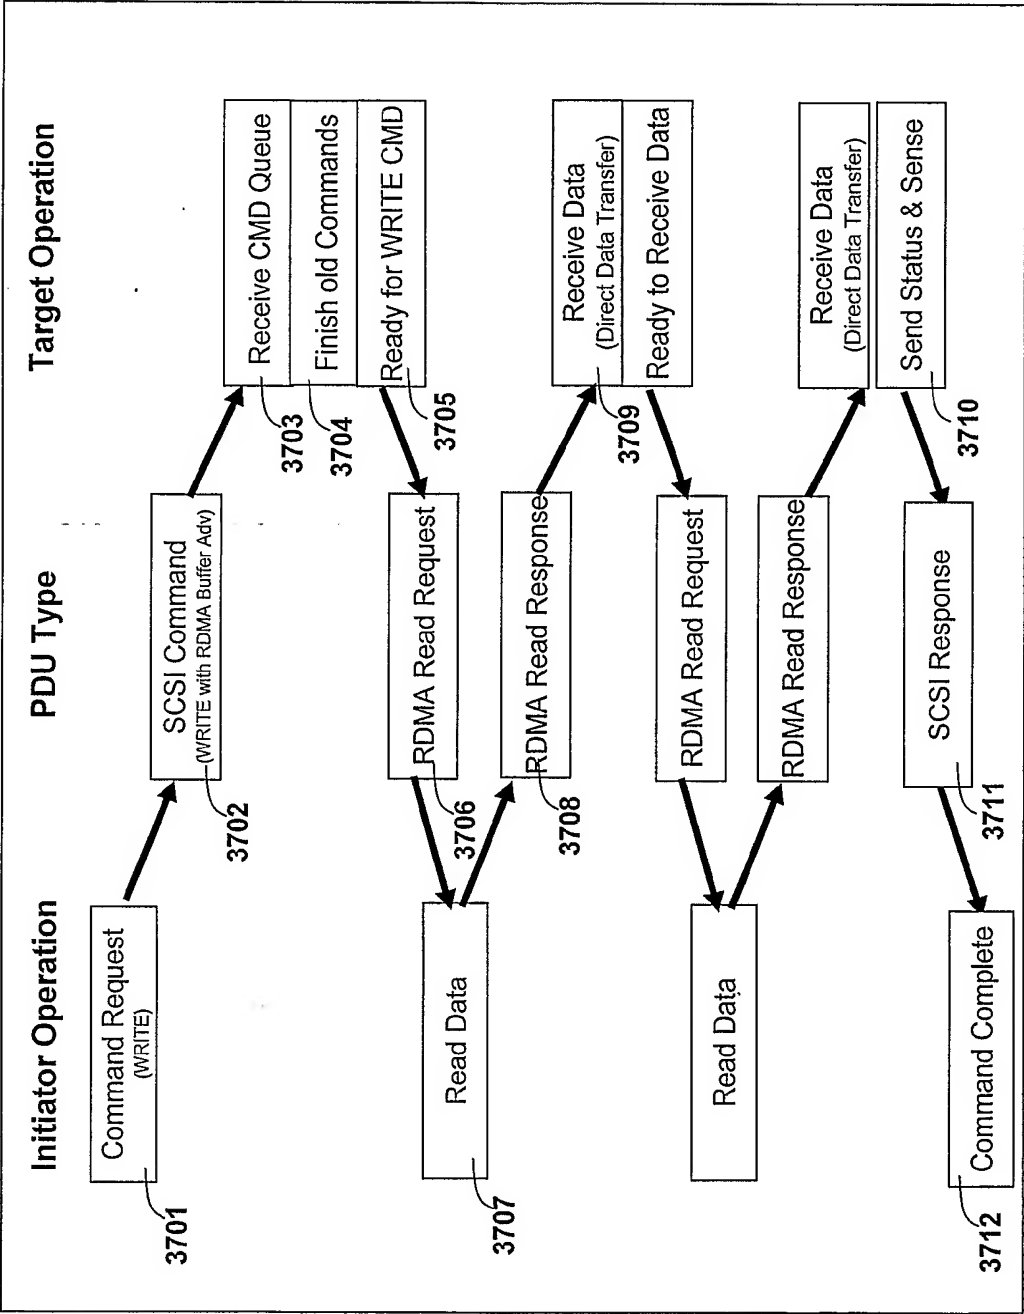
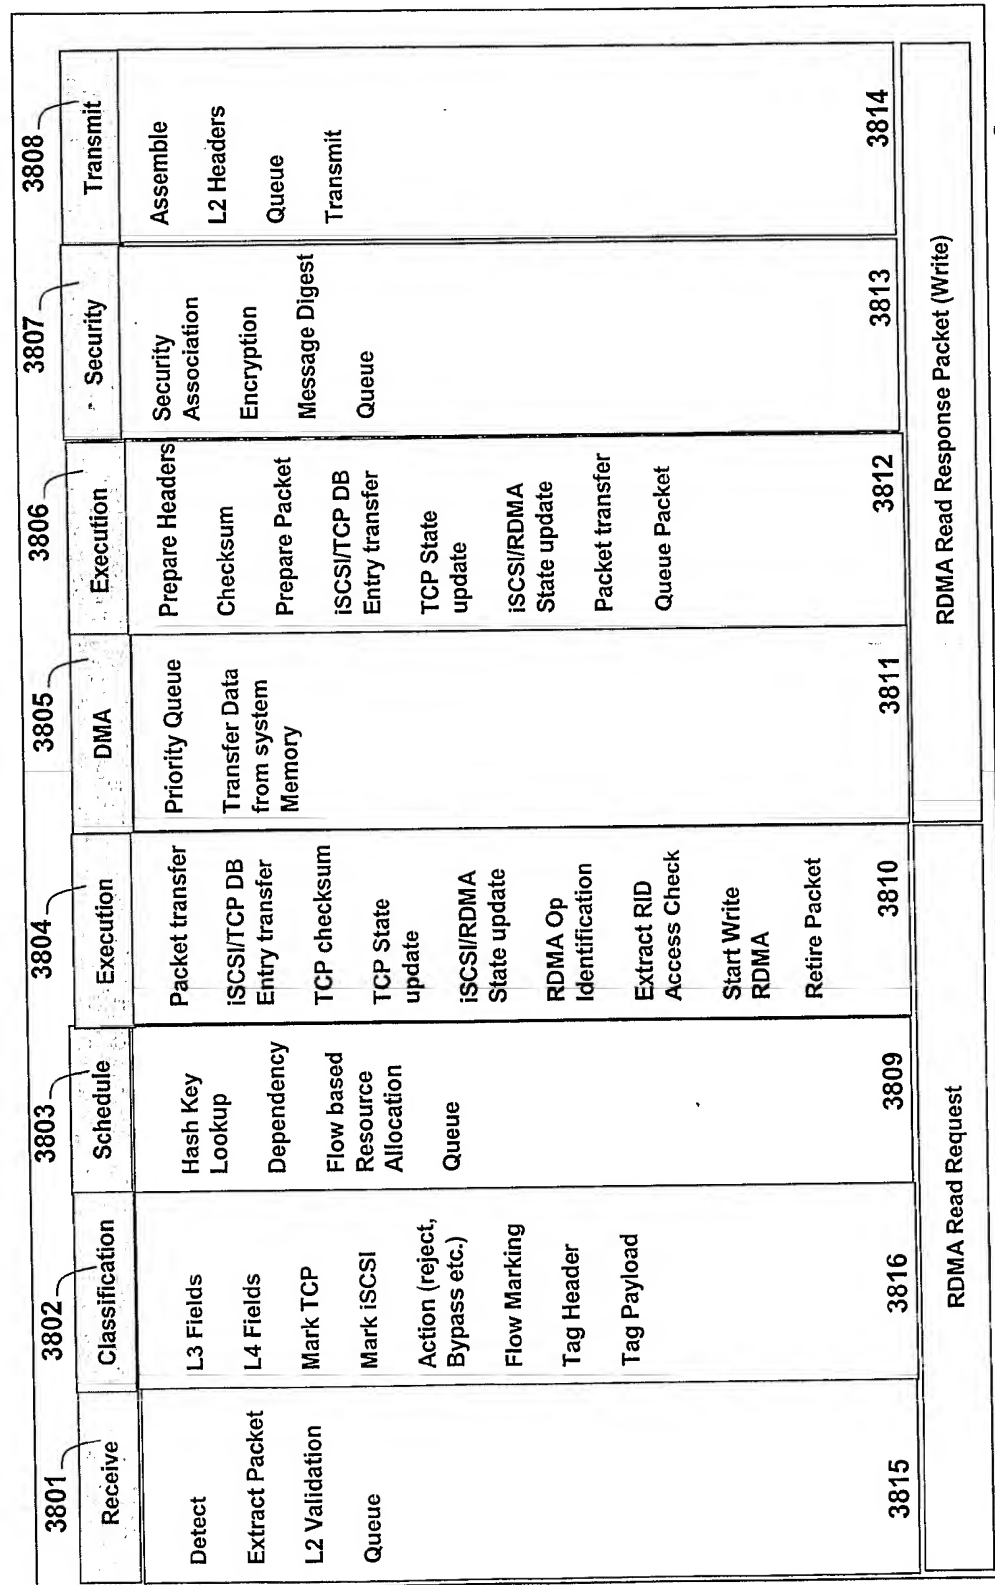


Fig. 38 Write Data Packet Flow (RDMA)



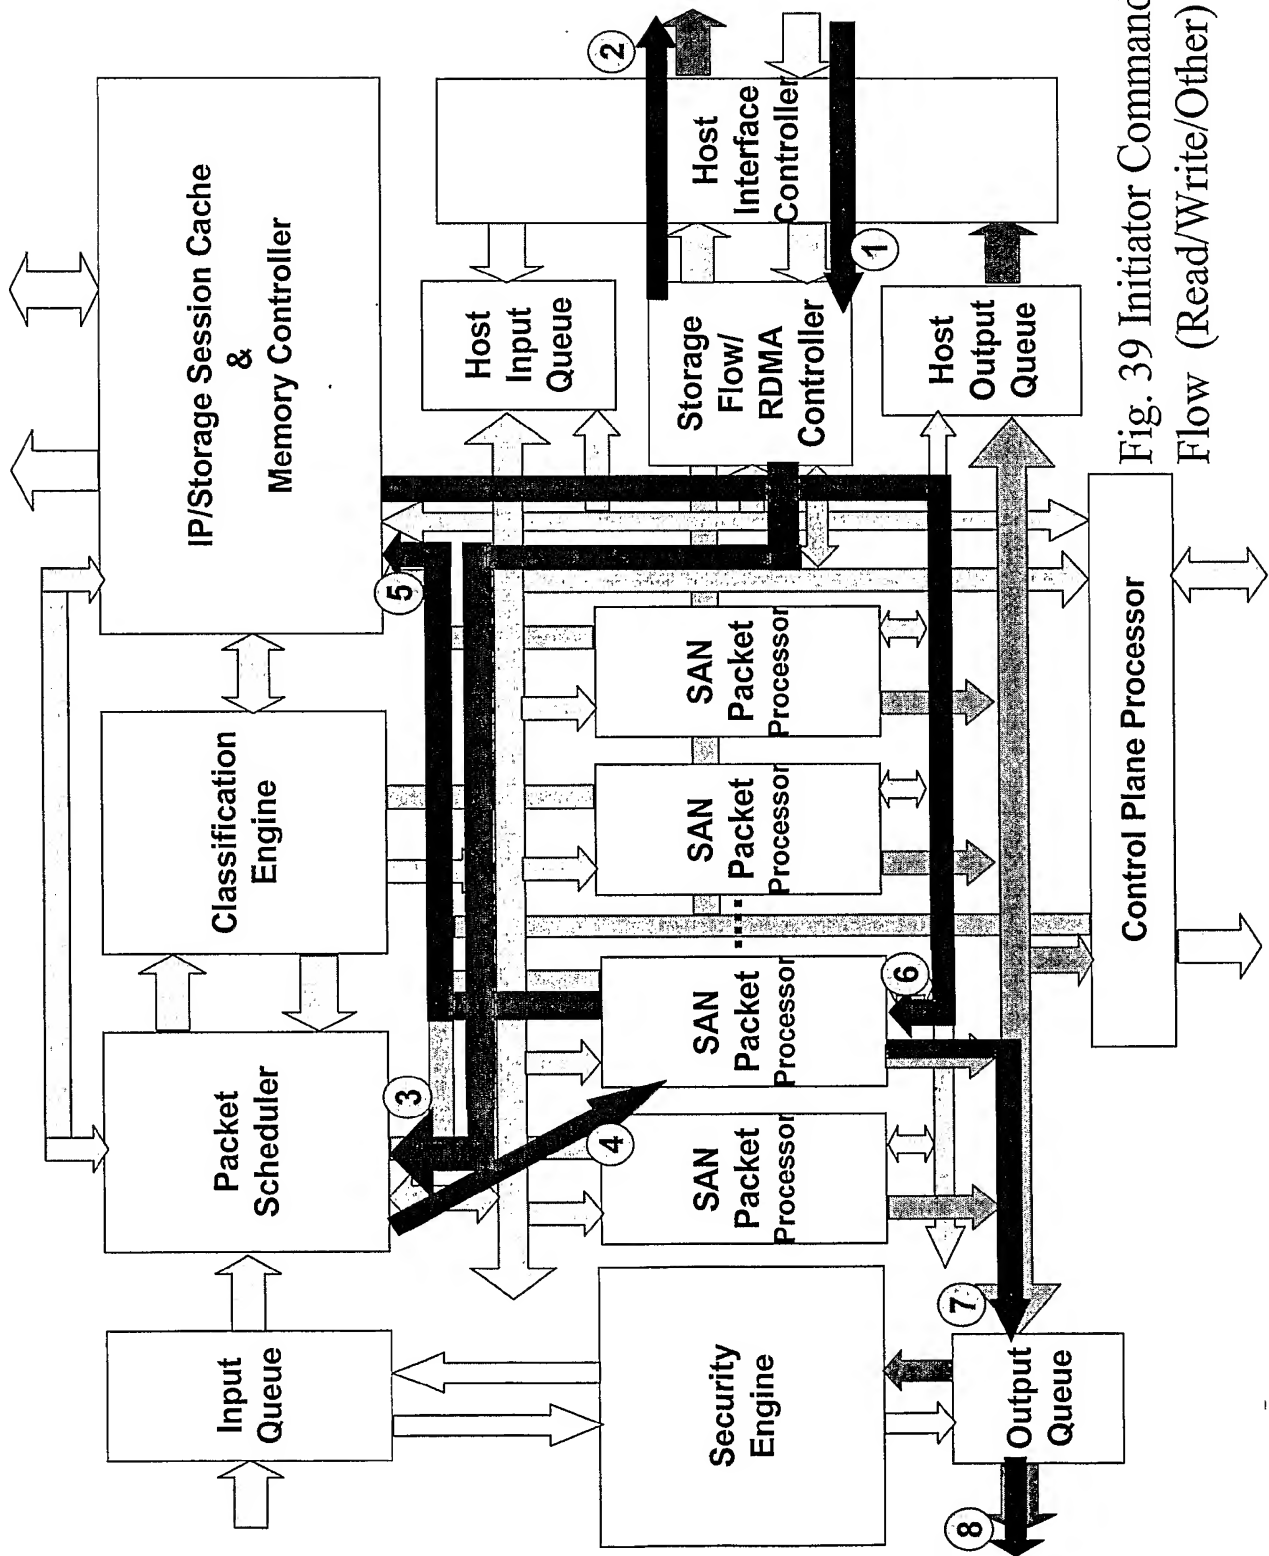


Fig. 39 Initiator Command Flow (Read/Write/Other)




Fig. 40 Read Data Flow

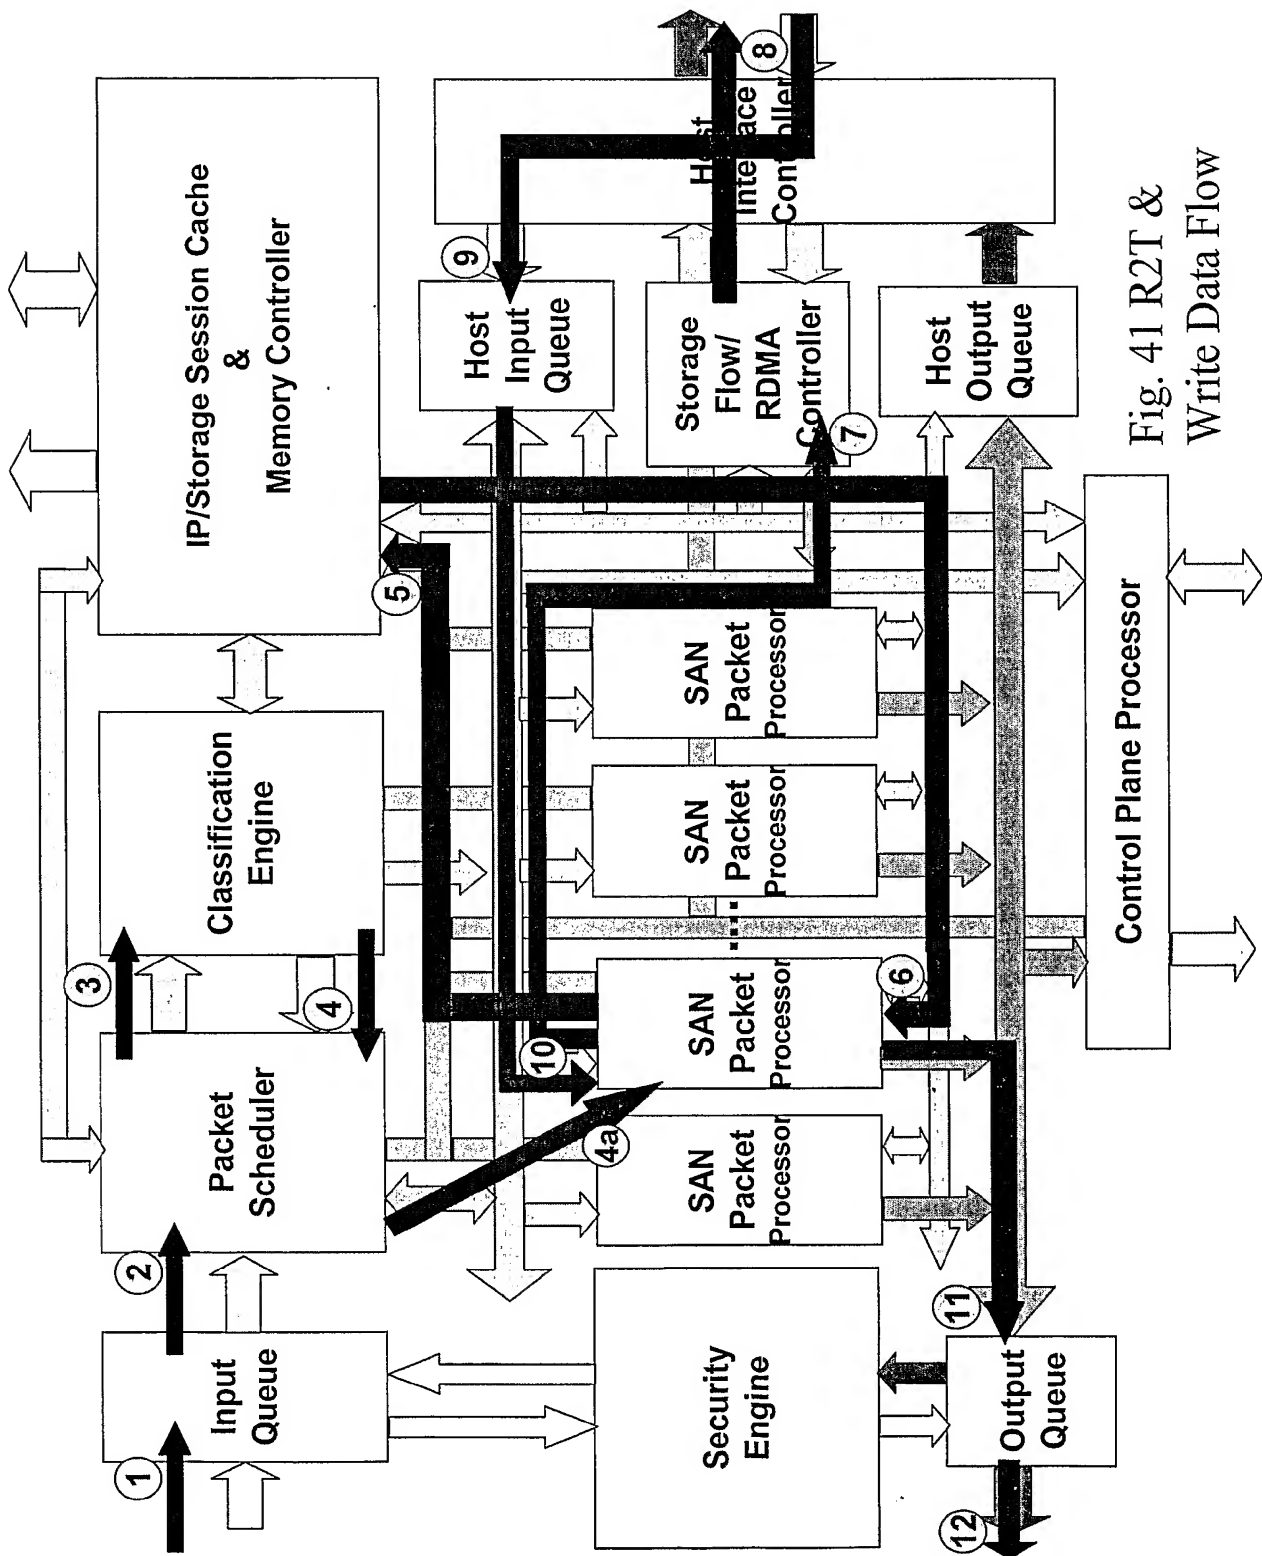


Fig. 41 R2T &
Write Data Flow




Fig. 42 Read Data Flow (Secure)




Fig. 43 R2T &
Write Data Flow (secure)

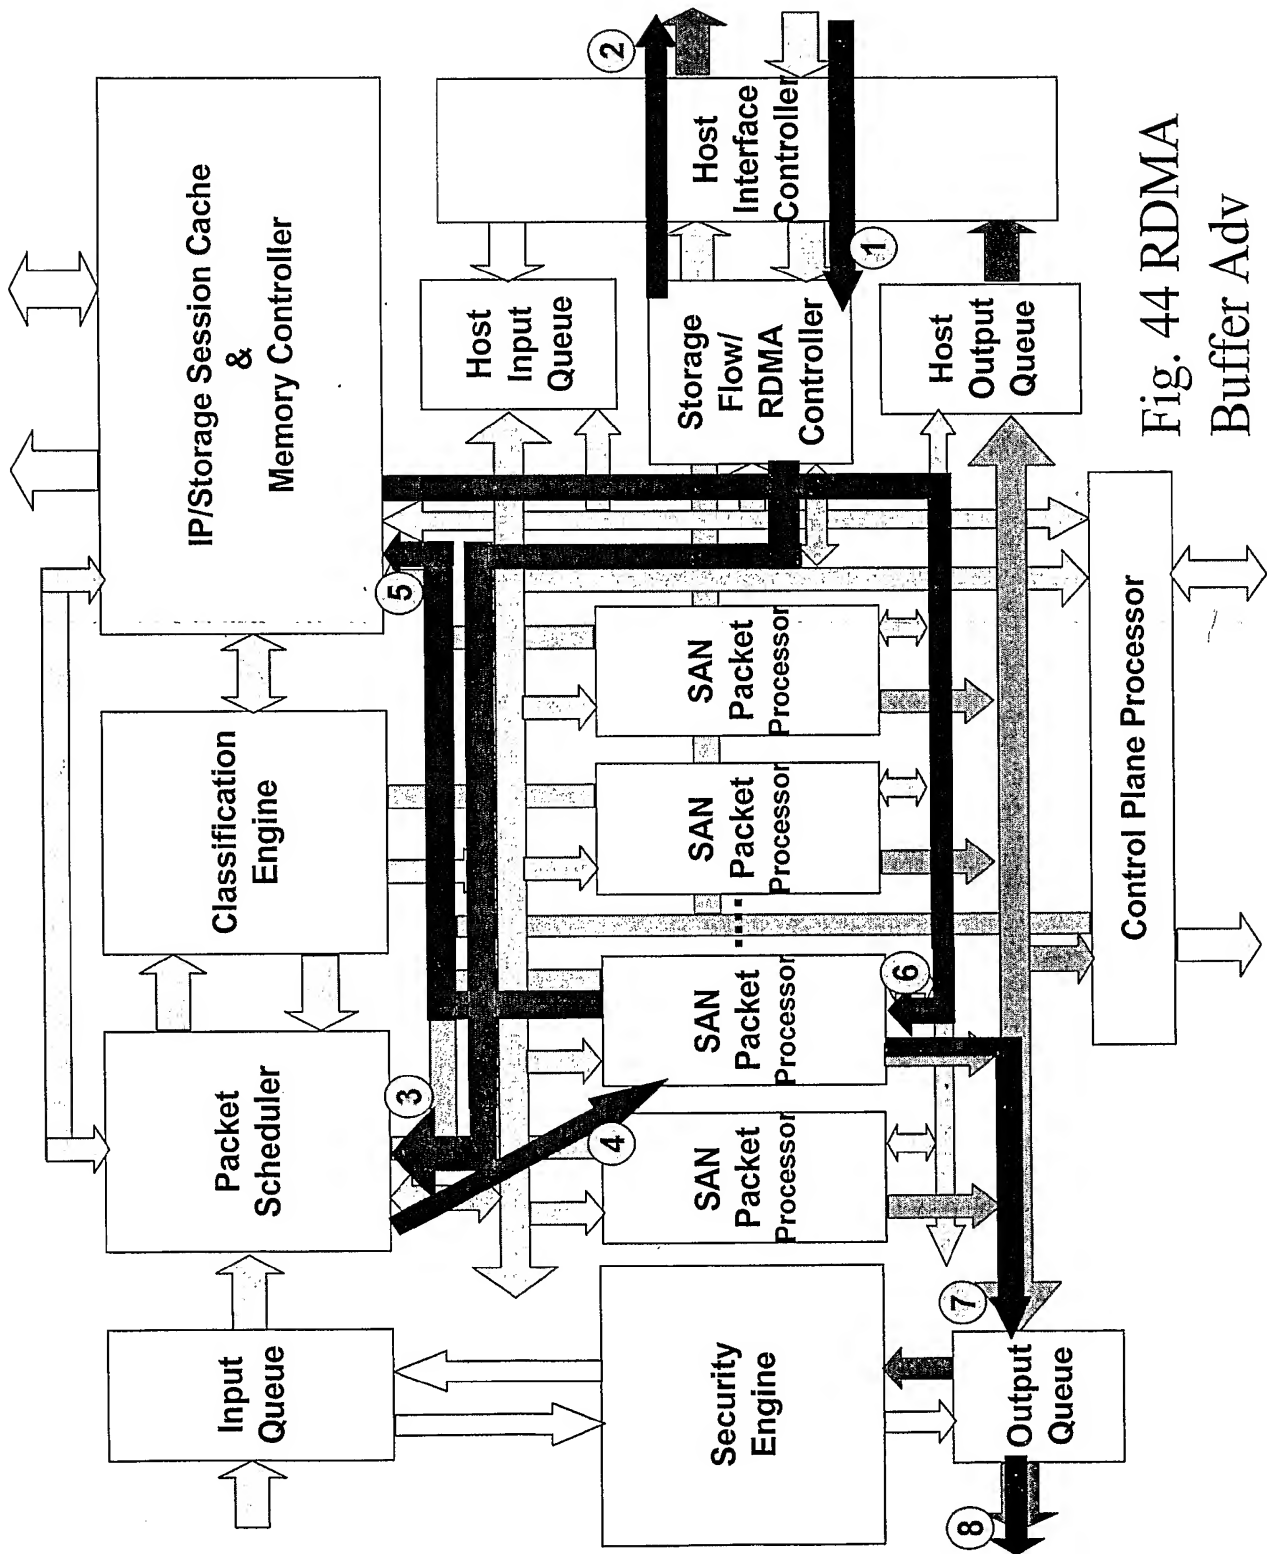


Fig. 44 RDMA
Buffer Adv




Fig. 45 RDMA Write Flow

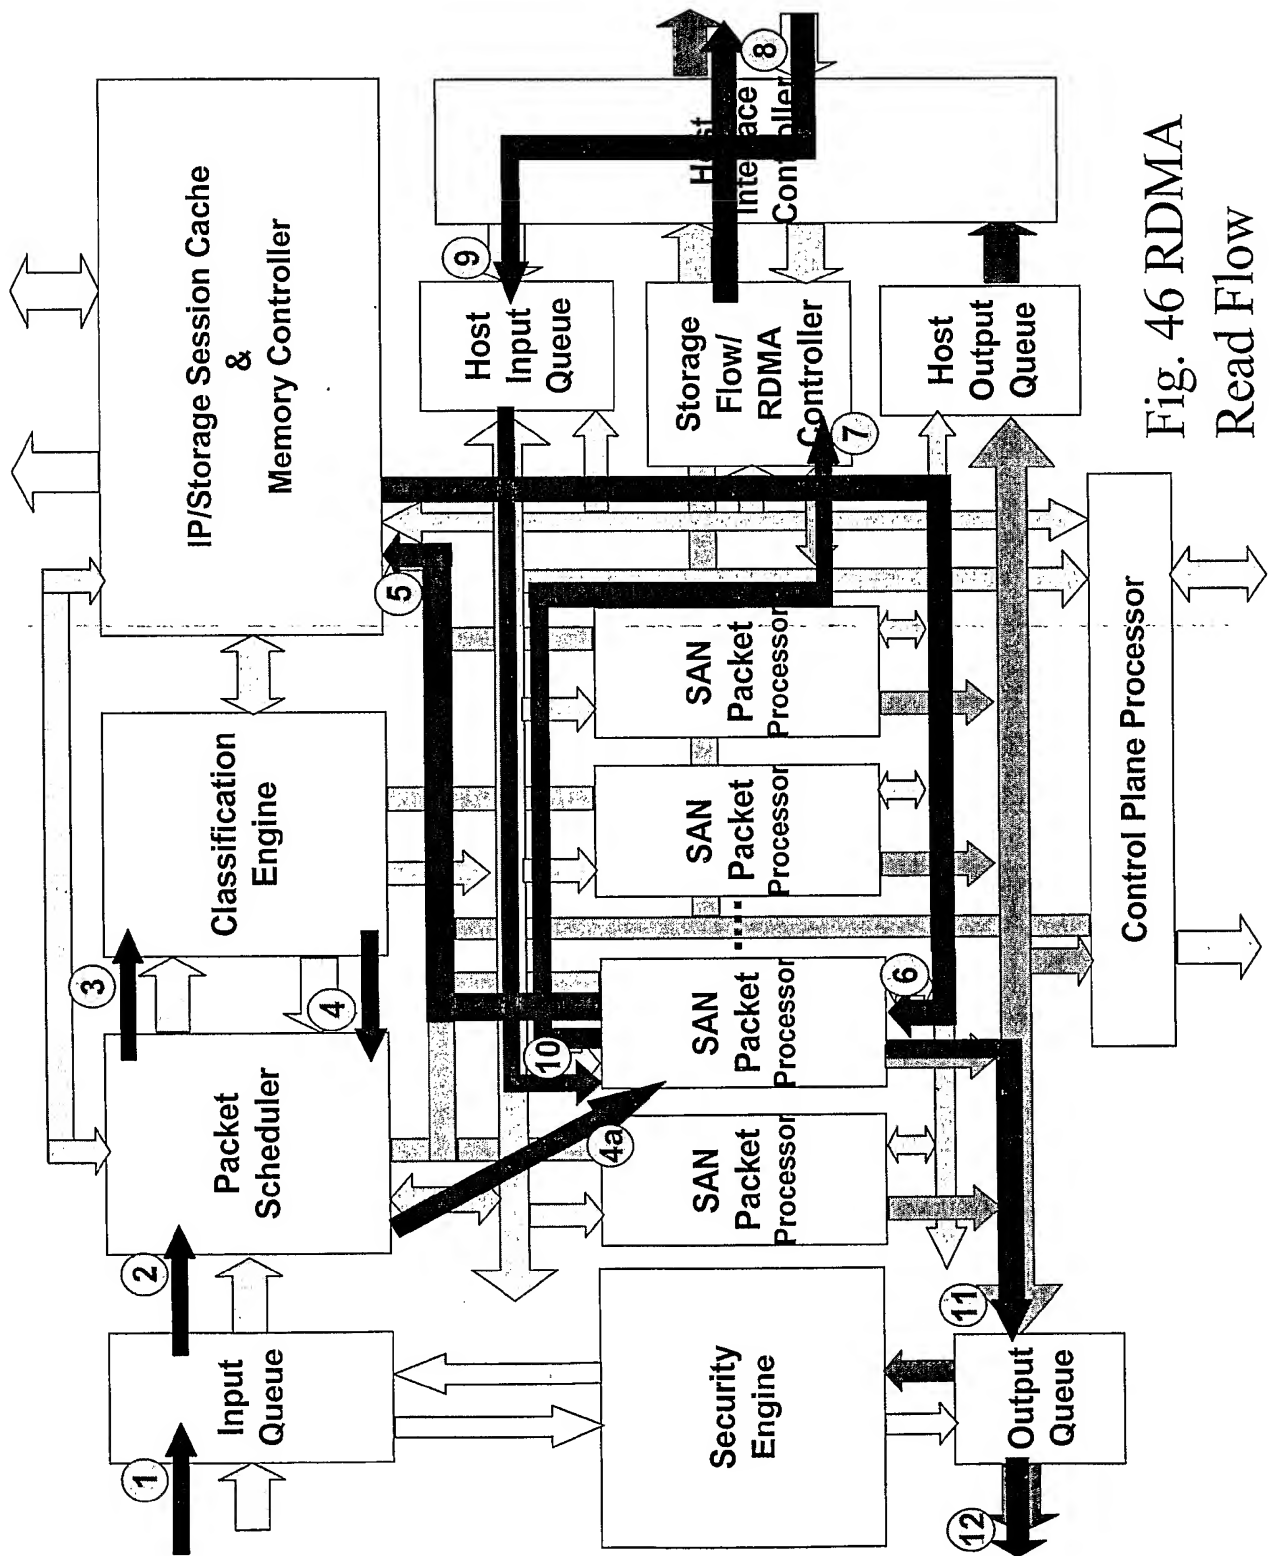


Fig. 46 RDMA
Read Flow





Fig. 47 Session
Creation Flow






Fig. 49 Target Session
Creation/
Tear down Flow




Fig.50 Target Write Flow (Send R2T).




Fig. 51 Target Write
Flow (Receive Write)

